

Try the latest **openSUSE 11.1** & **Sabayon 4**

Rs 125
ISSN 0974-1054

2 DVDs Free

LINUX

THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 06 ISSUE: 12 February 2009 116 PAGES

ISSUE# 73

ForYou



Still Stuck to Torrents?

Try Metalinks for downloads

Cross Swapping

for better multi-OS
PC performance

A-Z of Networking

for beginners

Get started with

GlassFish

application server

Build a highly-available

Web Server Cluster

Published by EFY—ISO 9001:2000 Certified



An exclusive interview with **Richard M. Stallman**

Leader of the Free World

**Life as an administrator
is complicated enough...**


**But Backup and Recovery for your
Linux Servers does not need to be.**

NetVault: Backup simplifies backup and recovery
without compromising **Functionality** and **Scalability**

NetVault: Backup provides unmatched Data Protection for all major variants of Linux. We are now offering you the chance to see just how good NetVault is at no cost.

We provide continuous data protection (CDP) for your Linux servers and advanced application protection and recovery for MySQL, PostgreSQL, Oracle, Exchange, DB2 to name just a few.

Permanent FREE Use Edition for Linux is available for download at <http://www.bakbone.com/nvbu/redhat/>



NetVault is a true Linux data protection solution featuring:

- ☐ Online backup
- ☐ Point and click recovery
- ☐ Fully automated protection and recovery support for Linux based applications
- ☐ Virtual Tape Library (VTL) with disk staging
- ☐ SAN support with LAN free backups.
- ☐ Backup duplication for off-site storage of backups
- ☐ Support for MySQL native replication, restore to table level and to alternative databases or instances.

BakBone
The Power of Simplicity™

For more information, please contact:

✉ : IndiaSales@bakbone.com
☎ : +91-11-42235156



Who said, "Opportunity
knocks only once" !!!

RHCE

at just **Rs. 2,500*** only

*Re-take exam
at
80% discount*

*Back pack
Pen Drive
Wrist Watch
RHCE Badge*

- 1) Enroll for Red Hat official curriculum and avail RHCE exam at Rs. 12,500/-
- 2) In case you do not pass your exam, retake the *RHCE exam at 80% discounted fee on standard RHCE exam Fees
- 3) Pass the RHCE exam and get RHCE goodies (back pack, pen drive, wrist watch and RHCE Badge)

Red Hat Certified Engineer (RHCE) is a performance based test that measures actual competency on live systems.



"RHCE called "The Crown Jewel of Linux Certifications" proves an individual's ability to configure networking services and security on servers running on Red Hat Enterprise Linux"

**OFFER
VALID ONLY TILL
FEB 15, 2009**

*(on 2nd and consecutive re-attempt(s))

* conditions apply



redhat

For more information visit www.redhat.in or email at training-in@redhat.com



An exclusive interview with RMS44



Leader of the Free World

FOR YOU & ME

- 18 The GIMP 2.6: Reaching Newer Heights
- 22 Distromania: A Multi-booting Template
- 25 Improve Multi-OS Computer Performance through Cross Swapping
- 28 Enabling Indian languages on the FOSS desktop—Part 3: Your Desktop, Your Language
- 30 Sabayon 4: Is It Ready For You?
- 34 Son... this is KDE 4.2 on openSUSE 11.1! Savvy?
- 38 FOSSTIVAL 2009: Hackers for Hackers
- 41 Close Encounters of the Free Kind!
- 44 Face to Face with RMS: An Interview
- 52 Recipes for Networking
- 60 Analyse Your Disk Space Availability
- 68 GeeP: The Geeks of Pune!



ADMIN

- 56 Are You RHCSS-certified?
- 82 Building a Highly-available Web Server Cluster

DEVELOPERS

- 88 What's in the Glass(Fish)? Part 2: Getting Started with the Application Server



CONTENTS

PLAYERS

- 64 “We have all that Red Hat has, but our prices will pretty much be half”
—An Interview with Shane Owenby, director, Linux & Open Source, Asia Pacific, Oracle

GEEKS

- 70 Building a Server from Scratch—Part 1: The Operating System
74 Metalinks for Download Anyone?
78 Programming in Python for Friends and Relatives— Part 10: Python on the Net Using CGI and WSGI

COLUMNS

- 76 FreedomYug: If You Love Someone, Set Them Free
96 CodeSport
99 The Joy of Programming: About the Java Overflow Bug
102 A Voyage To The Kernel: Segment: 2.2, Day 8



REGULAR FEATURES

- 06 Editorial
08 You Said It...
10 Technology News
16 Q&A Section
58 Industry News
87 Linux Jobs
94 Tips & Tricks
106 FOSS Yellow Pages



LFY DVD 1: openSUSE 11.1

The complete operating system that you can run across a range of desktops and servers.



LFY DVD 2: Sabayon4

If you are a gamer, or a multimedia and graphics buff, then this operating system is for you.

Note: All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported Licence a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence.

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)

D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 26810602, 26810603
Fax: 26817563
E-mail: info@efyindia.com

BANGALORE

No. 9, 17th Main,
1st Cross, HAL II Stage,
Indiranagar, Bangalore 560008
Ph: (080) 25260023; Fax: 25260394
E-mail: efyblr@efyindia.com

CHENNAI

M. Nackeran
DBS House, 31-A, Cathedral Garden Road
Near Palmgroove Hotel, Chennai 600034
Ph: 044-28275191; Mobile: 09962502404
E-mail: efychn2@efyindia.com

Customer Care

E-MAIL: support@efyindia.com

Back Issues

Kits 'n' Spares

D-88/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: (011) 32975879, 26371661-2
E-mail: kits@efyindia.com
Website: www.kitsnspares.com

Advertising

KOLKATA

D.C. Mehra
Ph: (033) 22294788
Telefax: 22650094
E-mail: efycal@efyindia.com
Mobile: 09432422932

MUMBAI

Flory D'Souza
Ph: (022) 24950047, 24928520; Fax: 24954278
E-mail: efymum@efyindia.com

PUNE

Zakir Shaikh
Mobile: 09372407753
E-mail: efypune@efyindia.com

HYDERABAD

P.S. Muralidharan
Ph: 09849962660
E-mail: efyhyd@efyindia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD

Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 24942538, 24925651,
24927383 Fax: 24950392
E-mail: info@ibhworld.com

Printed, published and owned by Ramesh Chopra. Printed at Ratna Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

Editorial

Dear Readers,

The buzzwords in 2009 seem to be 'recession', 'downturn', 'cost-cuts', and 'layoffs'. Naturally, the pressure on the tech community has also increased exponentially. Every discussion seems to start and end with "constraints on budgets". Customers (and our bosses) not only expect quality solutions to be delivered on time, but at lower costs too.

Can Open Source be our messiah during these challenging times?

Shane Owenby, director, Linux & Open Source, Asia Pacific, Oracle Corporation (interviewed on Page 64) surely thinks so. Many articles in leading business magazines, like *Fortune*, *eWeek*, etc, seem to suggest the same. Closer to home, many IT heads are suggesting that they are more interested today to understand the long-term financial and technical implications of Open Source solutions. But, they also outline a few challenges—many of which we have discussed earlier.

The most popular one continues to be "lack of support", or what is now being touted as "a visible lack of support". Just with the addition of a single word, "visible", the nature of the challenge that the open source community needs to overcome changes tremendously. What the decision-makers seem to be saying is that, despite the claims of adequate support for open source, they don't see much evidence of the same. If the proprietary world is looked upon as being monopolistic, there don't seem to be many choices in the open source world either. Hence, they prefer "a known devil to an unknown one".

The second issue that seems to bother most IT heads is that for every challenge there are just too many options in the open source arena! As one CIO comments, "For content management systems, we have Joomla, Drupal, Wordpress, and fifteen more options! How do I evaluate them all, and decide which one is suited for me?" But, aren't there hundreds of options in the proprietary world too? "Yes, there are, but the leading brands are well-known, and the eco-system has multiple points of references for IT guys to make an informed choice."

But aren't these two challenges in contradiction to each other—the first states that there's limitation of choice in the open source world, while the other claims that the abundance of options makes life confusing?

Not really! The main concern with open source remains that while there are multiple solutions for any IT challenge, the number of vendors who can support these solutions are "visibly" limited or absent.

So, what can we do? First, the vendors need to be more visible and get more aggressive. If they keep waiting for their organisation to grow into a multi-million dollar business before they start marketing or projecting themselves as a brand, they may be losing an opportunity to grow. Second, we need to create more platforms where competing solutions can be presented together and reviewed by decision-makers.

Can Open Source be our messiah during these challenging times?

LFY's FOSS Yellow Pages is an attempt to help young Indian organisations in the open source arena announce their services to decision-makers in the IT world, and Open Source India (www.osidays.com) in March this year shall attempt to create yet another platform for competing solutions to meet. But, if we want to accelerate the adoption of open source in India, we will need to come together and create many such platforms. Can we do that?

"Yes, we can!"

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@efyindia.com

Features

- Receive emails via SMS
- No GPRS required
- Use your existing Handset
- Works with any email service: corporate, POP3 and IMAP accounts.
- Compose, forward & reply immediately

Change your path, lead the fastest track!

Be the leader of business communication
with our innovative product. Feel the efficiency!

emergic
email2sms

Mail on Mobile

For Free Lifetime
email2sms service
login to
<http://m3m.in/register4free>

netCORE
The Innovation Company

sales@netcore.co.in | www.m3m.in

<u>MUMBAI</u>	<u>NEW DELHI</u>	<u>CHENNAI</u>	<u>BANGALORE</u>	<u>HYDERABAD</u>	<u>PUNE</u>	<u>AHMEDABAD</u>	<u>KOLKATA</u>
9820032324	9350101237	9380388685	9341609111	9392224001	9373727934	9824192328	9331420404



You said it...



Congratulations to the *LFY* team for making the New Year issue special. I am a 10+1 computer science student and I'm very interested in open source software. The articles and features about the Fedora 10 are excellent. As I'm interested in networking I have managed to get a load of information about the subject.

The CD team has done a very good job. All the software is very useful. The *Tips and Tricks* section is also very good. I have a request: could you try and include the Ubuntu Ultimate edition 2.0 and some graphical repositories? Hope the magazine improves by reducing spelling mistakes.

—**Anirudhan Adukkathayar,**
by mail

ED: Thanks for your feedback Anirudhan. We're glad that you liked the January issue and find the topics that were covered useful. Also, though the editing team does its best to avoid spelling and other factual errors, some escape our notice—we're all human after all. It'd be wonderful if you keep us informed about the errors you notice so that we can mention them in our 'Errata' section, if required.



I read a really cool article in the December 2008 issue of *LFY*—about how to create a portal with session cookies (cached on the client, and dynamic) and it helped me understand how cookies function. I have always had some doubts on how it works, but now I am very clear. I am trying to set up a portal on a Fedora 10 box, but have lost the magazine. Can you please let me know how I can get a copy of *LFY*'s December 2008 edition to continue working on my home set-up?

—**Gopal, by e-mail**

ED: The PDF of the article has been sent to you by e-mail. In case you want to order older issues you can always ping the chaps at kits@efyindia.com. :-)



I definitely enjoyed reading Sarath Lakshman's article titled, "udev Unplugged" in *LFY*'s Dec 2008 issue. Great work! I did want to point out, and ask about a couple of things:

1. A possibly useful addition: After one has edited/added rules into `/etc/udev/rules.d/<somefile.rules>`, the changes/additions generally don't take effect until the *udev* daemon is told to refresh its rules. This can be done by running the following command, as the root:

```
# udevcontrol --reload_rules
```

2. A query to the author: I don't know if it's just me, but the 'RUN' command does not seem to take effect on my Ubuntu box. I run Ubuntu 8.04.1 on the 2.6.24-22-generic kernel. The rule I added is as follows:

```
# Backup Disk: Transcend USB SATA 160GB disk.
SUBSYSTEM=="block",ATTR{dev}=="8:17",ATTR{st
tart}=="63",ATTR{size}=="312576642", RUN+="/
usr/bin/notify-send --urgency=normal --expire-
time=5000 'BACKUP DISK inserted (Transcend
160GB)...'",SYMLINK+="bkpdisk"
```

The thing is, the soft link `/dev/bkpdisk` is created, but the RUN command is not getting executed. (I also tried substituting the above command with an executable shell script, with the same result.) A very similar problem is alluded to at ubuntuforums.org/showthread.php?t=733615, where in the thread the author subsequently states that using full pathnames within the RUN command solved the problem. It does not seem to work with me. Any ideas?

—**Kaiwan N Billimoria, by mail**

Sarath replies: Regarding the first point, I don't think it is required. Because the latest versions of *udev* searched all the rules in `/etc/udev/rules.d/` for matches, every time. I just tried it now and verified this. Reloading the daemon is not really required. You can restart *udev* if required by running `/etc/init.d/udev`

reload' as the root. The *udevcontrol* command doesn't seem to be in Ubuntu 8.10.

Here is one problem with `RUN+="/usr/bin/notify-send`. We cannot call *notify-send* like that. Because, first of all, 'notify-send' is user-specific and requires the display argument (environment variable: DISPLAY); and second, it depends on the dbus-session, which is different for each user.

If you look at my article, you will find a link like the following one:

```
DISPLAY=:0 su $user -c 'notify-send "Backup
Complete";
```

Whatever rule you call from *udev* is run as root and its environment settings are different. We have to specify it explicitly. So I have specified the display setting as `DISPLAY=:0`. Moreover, since *dbus* is used, you have to specify the user name. So, you have to run *notify-send* as the specific user for which you want to display a notification. So it is better to use a script as `RUN+=argument`.

Modify your above code as:

Rule:

```
SUBSYSTEM=="block",ATTR{dev}=="8:17",ATTR{st
art}=="63",ATTR{size}=="312576642", RUN+="/usr/
bin/pendrive.sh",SYMLINK+="bkpdisk"
```

File: `/usr/bin/pendrive.sh`

```
#!/bin/bash
DISPLAY=:0 su username -c '/usr/bin/notify-send
--urgency=normal --expire-time=5000 'BACKUP
DISK inserted (Transcend 160GB)...'
```

It will work. You can verify this by attaching RUN with some other independent utility.

Thank you for your queries :-)

Please send your comments or suggestions to:

The Editor

LINUX FOR YOU Magazine
D-87/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: 011-26810601/02/03
Fax: 011-26817563
Email: lfyedit@efyindia.com
Website: www.openITis.com





Retain & Retrieve your mission critical emails for several years!!!

emergic

M@ilArchiva

Email Archiving & Compliance solution

- ✉ Compliance (SOX, HIPPA)
- ✉ Insurance against Litigation
- ✉ Policy, Security & Retention rules
- ✉ Indexing Content & Attachment
- ✉ Encrypted Backup
- ✉ Compression & De-Duplication

Key Benefits

- Seamlessly integrates with any mailing solution (Exchange, Domino, Linux)
- Automated Archiving without end user or administrator intervention
- Quick, Easy Search & Retrieval
- Back up system for lost emails
- No special database required

Partial Client List

- Geojit Financial Services Ltd.
- IDBI Capital Ltd.
- Apollo Sindhoori Capital Investment Ltd.
- Flat Products Equipment (I) Pvt. Ltd.
- IndusInd Bank Ltd

netCORE
The Innovation Company

sales@netcore.co.in | www.netcore.co.in/ema

<u>MUMBAI</u>	<u>NEW DELHI</u>	<u>CHENNAI</u>	<u>BANGALORE</u>	<u>HYDERABAD</u>	<u>PUNE</u>	<u>AHMEDABAD</u>	<u>KOLKATA</u>
9820032324	9350101237	9380388685	9341609111	9392224001	9373727934	9824192328	9331420404

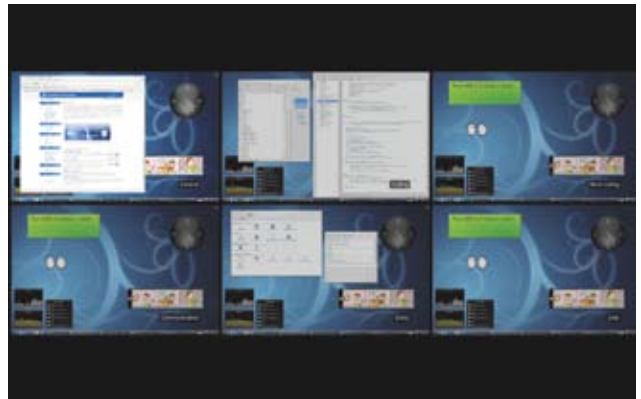
Technology News



KDE 4.2 knocks on your doors

By the time you read this, it will be at your doorstep as the final version is scheduled to go live on January 27. What we're looking at currently is the final testing version, KDE 4.2 RC. The new version introduces a large number of improvements over the v4.1.x branch. Some of the highlights according to the release information are as follows:

- **Compositing Desktop effects** are enabled where hardware and drivers support it, and automatic checks confirm that the compositing feature works before enabling it.
- **New desktop effects** have been added, such as the Magic Lamp, the Minimise effect, the Cube and the Sphere desktop switchers. All existing effects have been polished and feel natural due to the use of motion dynamics.
- **Central elements** of the desktop experience have seen significant improvements to give a more coherent experience. These include grouping and multiple row layout in the task bar, icon hiding in the system tray, notifications and job tracking by Plasma, the ability to have icons on the desktop again by using a Folder View as the desktop background.
- KRunner, the **'Run command' dialog box**, has extended functionality through several new plug-ins, including spell-checking, Konqueror browser history, power management control through PowerDevil, KDE Places, Recent Documents, and the ability to start specific sessions of the Kate editor, Konqueror and Konsole.
- The Plasma workspace can now load **Google Gadgets**.
- **Theming improvements:** A new System Settings



module—Desktop Theme Details, gives the user control over each element of various Plasma themes.

In addition to the overall usability improvements, the application-specific highlights include: Dolphin now supports previews of files in tooltips and has gained a slider to zoom in and out on file item views; Konqueror offers **increased loading speeds** by pre-fetching domain name data in KHTML. A find-as-you-type bar improves navigation in Web pages; KMail has a powerful and attractive **message header list**, and reworked attachment view; PowerDevil, the new KDE4 power management infrastructure brings in a modern, integrated tool for controlling various aspects of mobile devices; a **new printing configuration system** brings back a number of features that users have been missing in KDE 4.0 and 4.1; KRDC, the remote desktop client, improves support for Microsoft's Active Directory through LDAP; Kontact has gained a new **planner summary**, and support for drag and drop in the free/busy view.

Of course, this is in addition to all applications that have received bug fixes, feature additions, user interface improvements and generally more polish. We'd strongly recommend you keep an eye on www.kde.org and update to this release if you haven't already done so.



GenNext processor support with RHEL 5.3

Red Hat has released Red Hat Enterprise Linux 5.3. In the third update to version 5 of the enterprise Linux, customers will receive a wide range of enhancements like much increased virtualisation scalability, expanded hardware platform support and incorporation of OpenJDK Java technologies, claim Red Hat sources. Customers with a RHEL subscription

will receive the Red Hat Enterprise Linux 5.3 update, which is available for immediate download from Red Hat Network.

The primary new features of Red Hat Enterprise Linux 5.3 include: increased scalability of virtualised x86-64 environments, support for Intel Core i7 (Nehalem) processors, and the inclusion of OpenJDK. The update also includes enhancements covering many other

components, with the release notes documenting over 150 additions and updates. These new features and improvements range from the general kernel, device driver and architectural enhancements to simplified desktop networking and full support of the GFS2 filesystem. For more information on the features and updates, read the blog on www.press.redhat.com and www.redhat.com/rhel.

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 36 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

LLC102: Essentials of Linux OS
LLC103: Linux System & Network Administration
LLC203: Linux Advanced Administration
LLC303: Linux System & Network Monitoring Tools
LLC403: Qmail Server Administration
LLC404: Postfix Server Administration
LLC405: Linux Firewall Solutions
LLC406: OpenLDAP Server Administration
LLC408: Samba Server Administration
LLC409: DNS Administration

Courses for Developers

LLC104: Linux Internals & Programming Essentials
LLC105: Programming with Qt
LLC106: Device Driver Programming on Linux
LLC107: Network Programming on Linux
LLC108: Bash Shell Scripting Essentials
LLC109: CVS on Linux
LLC204: MySQL on Linux
LLC205: Programming with PHP
LLC206: Programming with Perl
LLC207: Programming with Python
LLC208: PostgreSQL on Linux
LLC501: Programming with OpenGL

RHCE Certification Training

RH033: Red hat Linux Essentials
RH133: Red Hat Linux System Administration
RH253: Red Hat Linux Networking & Security Administration
RH300/301: Red Hat Rapid Track Certification Course

RHCSS Certification Training

RHS333: Red Hat Enterprise Security: Network Services
RH423: Red Hat Enterprise Directory Services & Authentication
RHS429: Red Hat Enterprise SELinux Policy Administration

**New Courses
RH142 & RH184**



**RHCSS Exams in Bangalore
2, 3 & 4 March 2009
Register before 23 Feb 2009**


**RHCE Exams in Bangalore
on 6, 13 & 27 February 2009**

**RH301 from 9 & 23 Feb '09
RHS333 from 7 Feb `09
RH423 from 14 Feb `09
RHS429 from 21 Feb `09**

Linux Support & Solutions
Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS

For more info log on to:
www.linuxlearningcentre.com
Call: 9845057731 / 9343780054
Email: info@linuxlearningcentre.com

**RED HAT Training Partner
RHCE Examination Centre**



Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre

635, 6th Main Road, (Adj.. Bank of India) Hanumanthnagar, Bangalore 560019

Tel: +91.80.22428538 / 26600839 / 26610999 TelFax: +91.80.26600839

Cell: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com

URL: www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore

1291, 24th Cross, 30th Main,
BSK II Stage, Bangalore 560070

Tel: +91.80.26712928



Linux 2.6.28 is out with ext4 flagged

"It doesn't really matter what day it is, or what holiday (if any) you're celebrating, because even if you sit at home, alone in your dank basement, without any holidays or friends, I bring you tidings of great cheer: you can now download Linux-2.6.28, and compile it to your heart's content!" That was Linus Torvalds on Christmas Eve last year. "Listen to the cheerful grinding of your hard disk as you reboot into an all-new kernel—and I'm sure that if your computer could smile, it would have a big silly grin on its non-existent face. So as you sit there in your basement, give your computer the holiday cheer too."

Among the prime highlights of this release is obviously the ext4 filesystem being tagged stable. The backwards-compatible replacement of ext3 introduces bigger filesystems and file sizes, extents, delayed allocation, multi-block allocation, improved block allocation algorithms, faster *fsck*, online defragmentation, and faster and more robust journaling. Along with this release debuts something called, "Graphic Execution Manager" (GPM), a memory manager for the GPU memory. The GPM works as a central manager for buffer object placement, caching, mapping and synchronising.

Other features offered by the new kernel include support for Ultra Wide Band (UWB), wireless USB and UWB-IP; memory management scalability improvements; and container freezer, a cgroup subsystem that utilises the swsusp freezer to freeze and restart an arbitrary group of tasks determined by the user. Then there is a boot tracer to help developers optimise boot times; disk shock protection; along with the inclusion of the staging branch that contains incomplete or unstable drivers. For details of all the features, refer to kernelnewbies.org/Linux_2_6_28. Note that an update to the new kernel version has been made available with a wide range of bug fixes on January 18. Linux 2.6.28.1 is available for download from www.kernel.org


Amarok upgrades to 2.0.1.1 with some handy features

Amarok 2.0.1.1 was released on January 11. This release brings back some of the features a lot of us have been waiting for, and which were absent in v2.0. Those of you who were used to the queuing, playlist search and filtering, as well as 'stop after current track' features from the Amarok 1.4 series, and were griping because of their absence when v2.0 came out, have no reason to make a face any more—these features are back again. In addition, the long-awaited sorting of the collection based on the composer has also been introduced in this release.

The developers have also "...extended the Scriptable Service API to allow smoother integration of service scripts. The LibriVox service script has received some love and comes with new icons amongst other improvements. Media device handling has also seen many improvements: MTP devices can now delete multiple tracks at once, and the status bar gives visual feedback when deleting tracks from iPods." Of



course, many of the annoying bugs that were identified in the previous release have also been quashed, like "...the collection scanner now doesn't miss entries anymore and completes its work, and the user now gets a warning when the scanning fails." Download the latest release from amarok.kde.org



Real-Time Messaging Protocol to be opened up

Adobe has announced plans to publish the Real-Time Messaging Protocol (RTMP) specification, which is designed for high-performance transmission of audio, video, and data between Adobe Flash Platform technologies. Providing developers and companies open and free access to RTMP is the latest advancement of the Open Screen Project. This is an industry-wide initiative to enable the delivery of rich multi-screen experiences built on a consistent runtime environment for Web browsing and standalone applications across PCs, mobile devices, and consumer electronics.

RTMP provides an enhanced and efficient way to deliver rich content. Developers and companies will have free and open access to the documented RTMP specification to help enable unparalleled delivery of video, audio and data in the open AMF, SWF, FLV and F4V formats compatible with Adobe Flash Player. The RTMP specification is expected to be posted on the Adobe Developer Connection [www.adobe.com/devnet/rtmp] in the first half of 2009.

Sidux 2008-04, based on Debian Sid as of 2008-12-22

Although the release date of Debian Lenny is still uncertain, there's good news for its users. The Sidux developers have released version 2008-04—a full-featured Debian Sid-based live CD with a special focus on hard disk installations, a clean upgrade path within Sid and additional hardware and software support, enriched and stabilised with Sidux' own packages and scripts. The specific package versions for this release include: Linux 2.6.27.10 (smp, hard preemption), X.org 7.3, KDE 3.5.10, support for newer Intel chipsets (P4x, G4x and Q4x), OpenJDK 6, KVM support, etc. Visit sidux.com/Article303.html to download the ISO.

Android support now on Snapdragon chipsets

Wind River has developed a software code that will enable support for Android mobile software running on larger screen-sized mobile computing devices using Snapdragon chipsets from Qualcomm. This combination, according to the companies, will help create opportunities for mobile manufacturers to more quickly bring to market a wide range of Android-based WVGA devices. The software code will be contributed by Wind River to the Open Handset Alliance (OHA) repository in the near future.

SMART Notebook 10 for core Linux distributions

SMART Technologies has announced the availability of SMART Notebook 10 software for the latest core distributions of Linux. SMART Notebook 10 software now supports five core Linux distributions, including Ubuntu 8.04, openSUSE 10.03, Red Hat Enterprise Linux 5.1, Fedora Core 9, and Debian 4.0 (Etch) r3.

The SMART Notebook 10 software has new features. For instance, with Table Tool, you can easily create a table from the SMART Notebook toolbar, and then insert, or drag and drop text, images and objects into any cell, as well as add or delete individual cells to create asymmetric tables. There is Object Animation to animate any SMART Notebook object with effects such as fading in, flying in or spinning. Magic Pen is a tool used to spotlight, magnify, or zoom in on an image, or write notes that will disappear in 10 seconds. Active Alignment helps to format a SMART Notebook page to align objects to other objects; and Shape Pen is to draw a freehand shape on the SMART Board interactive whiteboard, while SMART Notebook recognises and perfects it. For details visit www2.smarttech.com/st/en-US/Products/SMART+Board+software

Check out ULE scheduler and DTrace on FreeBSD 7.1

FreeBSD has announced the second release from the 7-Stable branch, which improves on the functionality of FreeBSD 7.0 and introduces some new features.



Some of the highlights are that the ULE scheduler is now the default in generic kernels for AMD 64 and i386 architectures, which is supposed to significantly improve performance on multi-core systems for many workloads; support for using Dtrace (a comprehensive dynamic tracing framework) inside the kernel that has been imported from OpenSolaris; and a new improved NFS Lock Manager (NLM) client.

Other additions include: boot loader changes that allow booting from USB and GPT-labelled devices; addition of the `cpuset(2)` system call and `cpuset(1)` commands, providing an API for thread to CPU binding, and CPU resource grouping and assignment; KDE updated to 3.5.10; and GNOME updated to 2.22.3. For a complete list of new features and known problems, check the release notes and errata list at www.FreeBSD.org/releases/7.1R/relnotes.html and www.FreeBSD.org/releases/7.1R/errata.html, respectively. The update is available for the AMD 64, i386, IA64, PC98, PowerPC, and Sparc64 architectures. You can grab an ISO from [ftp://ftp.freebsd.org/pub/FreeBSD/](http://ftp.freebsd.org/pub/FreeBSD/)

Storage gets united and flexible with Tyrone Opslag FS2!

Netweb Technologies, a server, storage and HPC solution provider, has launched its new product—the Tyrone Opslag FS2 series of storage solutions. Unifying the traditional Storage Area Network (SAS) and Network Attached Storage (NAS) into a single box, FS2 offers higher flexibility and allows users to have as much of both. The product provides scalability from a few terabytes to 576 TB, and could perform up to 2GB/s.

Apart from NAS and SAN functionality, Opslag FS2 can also function as a VTL (Virtual Tape Library). It helps users shift back-up operations from the slow tape drives to faster disk-based storage, without causing any major changes to the existing set-up. FS2 runs on Quad-Core Xeon processors and supports up to 16 GB of memory. It supports various redundant arrays of independent disks or RAIDs, which ensures protection of valuable data even in case of multiple hard disk drive failures. The RAID cache is protected by a battery back-up unit, and the system also supports multiple snapshots with scheduling and volume/share replication. For customers who do not wish to take any chances, FS2 supports block-level replication to a second similar unit along with failover/high-availability.

The most important and innovative feature in FS2 is its support for SRP (SCSI RDMA Protocol). This protocol helps to support block level access using an Infiniband adapter and delivers high-sustained throughput. Visit www.tyronesystems.com/fs2 for more information.





India's leading Open Source event is coming...

OPEN SOURCE INDIA: TECH DAYS

It's happening at...

CHENNAI

It's happening on...

12, 13 & 14 of MARCH 2009

It will be useful for...

SOFTWARE DEVELOPERS AND IT MANAGERS
(And Newbies too!)

And, we seek YOUR help to make it the best in Asia.

Shoot us an email at countmein@osiweek.com, and let us know how you can help us in building
an Indian Open Source event into Asia's best.

Sponsors: _____



Powered By: _____



Vision Statement of Open Source India

To build a platform where-by we can enable India, and Asia, to lead in the development and adoption of Open Source software solutions.

Mission of OSI Tech Days 2009

- Development of Open Source software
- Promote usage of Open Source platforms and tools for software development
- Update software developers with the latest trends, technologies and best practices from the open source world
- Adoption of Open Source software
- Accelerate implementation of Open Source powered IT solutions
- Equip IT managers with latest open source solutions and best practices

Key Industry Verticals Being Targeted

- IT Firms (Software Development)
- ITeS Firms (BPO, KPO, etc)
- Banking & Finance (BFSI)
- Manufacturing
- Research and Development Labs
- Education (Government and Private Institutions)
- Government

Size of firms being Targeted

- SME (Small & Medium Enterprise)
- Enterprise

Schedule of Events

Three days plus four halls allow us to try and discuss many topics and reach out to a variety of audience. Here's a list of events that have been finalized till date:

Software Developers' Track

- **12th March:** FOSS-enabled software development
- **13th March:** FOSS-powered Mobile phones (and devices)

IT Managers' Track

- **12th March:** FOSS Tools for IT Managers
- **13th March:** Cloud Computing, Virtualisation and SaaS

.INFO Track

- **12th March:** FOSS-powered eGovernance
- **13th March:** FOSS for education (Half-day)
- **14th March:** FOSS-powered desktops

Other Events

- **12th March 2009:** IOSN Asia Summit
- **12th March:** Industry Demos (Half-day)
- **13th March:** Industry Demos (Half-day)

Media Partners: _____



Exhibitors and Sponsors may contact us at: sales@osiweek.com



KNOW HOW



Q I have installed a Linux machine without GUI. Please let me know how can I set the time on the command prompt? I can only log in to it by using SSH.

—Kumar Srinivas, by email

A. You can use the *date* command to set the time if you are not using NTP or you are not connected to the Internet. You can set *date* and *time* manually using the following command from the shell prompt as root:

```
date -s "22 JAN 2009 20:10:00"
```

or:

```
date +%Y%m%d -s "20090122"
```

Or, to just set the time, you can use:

```
date +%T -s "20:13:10"
```

...where 20 is hours, 13 is minutes and 10 is seconds. You can use *%p* if you want to set time in AM or PM format.

Q I have Fedora Core 6 installed on my machine. I was running out of space due to the low capacity of the hard disk drive. I have bought a new hard drive of 160 GB, but am unable to understand how to install and format it. I have physically connected the HDD as primary slave and the same is also detected by my system BIOS. How do I format

and use this hard disk?

—Priyanka Trivedi, by email

Since your new hardware is recognised by the BIOS, you need to check through your OS the device name that has been assigned to it. Run the following command as root to know the details:

```
# ls /dev/hd*
```

or:

```
# fdisk -l
```

This will give you the device name that your new hard disk has been assigned (in your case it should be *hdb*).

After you know the device name you need to divide the hard disk into partitions. To do this you can use the *fdisk* command as follows:

```
# fdisk /dev/hdb
Command (m for help):
```

You can press *m* to get the list of options. Press *p* to get the listing of the partitions—in your case this should display no partition, as it's a new disk.

So, let's create a new partition by pressing *n*, and follow the on-screen display:

```
command (m for help): n
command action
e   extended
p   primary partition (1-4)
```

As you can see from the terminal output, you can choose *p* to make a primary partition. *fdisk* always sets partition to 83 (Linux type) by default, although you always have options to play with different types of filesystems. Once you are done with partitioning, just press *w* to

write the modifications permanently.

Next, format the partition you have just created. To format the new disk you can use the *mkfs* tool. This tool is used to build a Linux filesystem on a device—usually a hard disk partition (assuming that you have created only one partition):

```
# mkfs -t ext3 /dev/hdb1
```

This will format your hard disk as an *ext3* filesystem. Once *mkfs* has finished its work, your partition is ready to be mounted and used. To mount the new partition, just create a folder in the */media* directory named *hdb1* or anything you wish, and run the following command:


```
# mount -t ext3 /dev/hdb1 /media/hdb1
```

This will mount your partition. However, the next time you reboot your system, you have to manually *mount* it again. To mount your partition automatically every time you boot your system, you need to append the following line in the */etc/fstab* file:

```
/dev/hdb1 /mnt/hdb1 ext3 default 2 1
```

Q I have an old Linux with Openoffice 2.0 installed on my laptop. I have few true-type fonts that I used with Windows. Is there any way by which I can use these fonts in my OpenOffice.org too?

—Sebastian John, by email

I'm assuming you want to use these fonts in your Fedora system. In that case, simply copy the fonts in your */usr/share/fonts/* directory—remember, you need root privileges for this action. You should now be able to use these fonts in OpenOffice.org and other assorted Linux applications. **END** 

17th Convergence India 2009

International Exhibition & Conference

Pragati Maidan, New Delhi, India

18-20 March 2009

*South Asia's largest **ICT** Event*
*Digital Convergence changing the **ICT** Landscape*



Featuring

- Telecom
- Mobility
- Information Technology
- Information Security
- Broadcast
- Cable
- Satellite

Certified by



Department of Telecommunications
Ministry of Communications & Information Technology
Government of India



Supported by



Department of Information Technology
Ministry of Communications & Information Technology
Government of India



Ministry of Information & Broadcasting
Government of India



Supporting Journal



Organised by



Exhibitions India Pvt. Ltd.
(An ISO 9001:2000 Certified Company)

217-B, (2nd Floor) Okhla Industrial Estate, Phase III, New Delhi 110 020, India Tel: +91 11 4279 5000 Fax: +91 11 4279 5098/99

Bunny Sidhu, Vice President, (M) +91 98733 43925 bunnys@eigroup.in /
Sambit Mund, Group Manager, (M) +91 93126 55071; sambitm@eigroup.in

Branches: Bangalore, Chennai, Hyderabad, Mumbai, Ahmedabad, California

www.convergenceindia.org

The GIMP 2.6

Reaching Newer Heights

Getting your hands dirty with the latest version.

*T*he GIMP, or the GNU Image Manipulation Program, is one of the most powerful free and open source raster image editing software. It rose from the ashes of a hideously crafted class project under the name of General Image Manipulation Program by Spencer Kimball and Peter Mattis. Later on, it was merged into the GNU Project and is often considered a replacement for the icy-pricey Adobe Photoshop.

Until recently, the diverse GIMP was

very much a mixed bag—too perplexing for newbies on the one hand while requiring high-end colour channels and plug-ins support. Besides, its meagre 8-bit colour support (instead of 32-bit) and complex UI was a major barrier in its adoption among professionals.

Finally, after eight years of planning, the developers took a major leap in GIMP development. The spanking new GIMP 2.6 was introduced, loaded with GEGL, a CMYX/32-bit colour channel, minor UI changes and a plethora of nifty add-ons and plug-ins.

On October 1, 2008, the venerable raster image editor got refurbished with a new version 2.6.x. The major change in this version lies in the core of the software—the addition of GEGL, a powerful graph-based image-editing framework.

Let's take an in-depth look at what the newer GIMP has to offer. In this review, I will be using the third revision of the latest GIMP—that is, GIMP 2.6.3, which provides significant improvements over the basic release.

When initiating the GIMP, you will be

GEGL and CMYK integration

GEGL, or Generic Graphics Library, is a powerful graph-based image-editing framework. The most enticing reason for developers to opt for GEGL is its library functions, which can deal with the 32-bit colour channel. In addition to this, GEGL supports the CMYK (cyan, magenta, yellow and 'k' for black) colour model, which is used in high quality colour printing.

Previous versions of the GIMP had support only for 8-bit colour channels, which was a setback for raw formats, and thus owners of high-end digital cameras were forced to opt for paid alternatives. The newer version paved the way for GIMP's acceptance and future development.



greeted with the latest splash screen—not the best, but certainly much better than its previous version. I always felt that if a bigger and wider splash were added by default, like Photoshop [PS], it would provide a more aesthetic feel. Unlike other image editing applications, GIMP's hardware resource expectations aren't that high, and likewise the application is quite responsive despite being one of the most feature rich editors.

The first thing you will notice about GIMP is its rejuvenated interface. Yes, it has finally gone through some UI changes. A new permanent 'empty' window will appear next to the toolbox and layer box. In addition to this, the toolbox, layer box, and active image editing window list has been merged into one window list, making GIMP more aesthetically appealing.

Apart from that, the current UI changes are too minimal and experimental—it seems developers were uncertain about what to implement and what not to. The current UI development indicates that developers intend to evoke community feedback for further UI development.

On the templates front, GIMP hasn't added new window sizes and ratios, and it would have been a welcome change if the developers had added the HD template by default.

In a nutshell, it's the tools, plug-ins and the core that have witnessed a major facelift. With the inception of GEGL, a few new image editing plug-ins and enhanced editing tools are now available.

We'll take a look at some of these new additions.

Polygon selection tool

The most exquisite feature of GIMP is the polygon selection tool, an add-on to the existing free-hand selection tool. Unlike the earlier versions, this tool lets you select/outline an image or its part in multiple clicks. Previously, you couldn't perform this function completely if you took your finger off the mouse, mid-click. So, in order to select/outline a part, you had to complete it at one go, else it wouldn't be done.

Brush dynamics

Brush dynamics is another value-added plug-in. The latest version now has added extensions to the brush tool—viz., Velocity, Pressure and Random. You can create different styles using similar brushes by varying the speed, although this option is aimed at users owning a pen and tablet. Each of the above mentioned parameters can be further enhanced by opacity, hardness, size and colour.

An enhanced text tool

The new text tool lets you frame your text into a box, and then you can move/resize it at will. This function is useful for large text editing, though you have to edit the text in a different window. On-canvas text editing has been planned for GIMP 2.8.



Figure 1: The new interface




Figure 2: Polygon free selection tool

The ability to pan beyond the border

This feature lets you use the brush tool at the edge of the image while the image is being zoomed in. The feature provides precise control over the brush tool making it much easier for you to control even at high zooms.

GEGL operations and tools

The talk of the town is, however, GIMP's GEGL integration. As already mentioned, GEGL supports high-end graph-based editing with 32-bit colour channel support. And although version 2.6 carries the GEGL tag with honour, is GEGL integration really worth it? No! There are many reasons GEGL doesn't make version 2.6 much better than 2.4 because the developers haven't

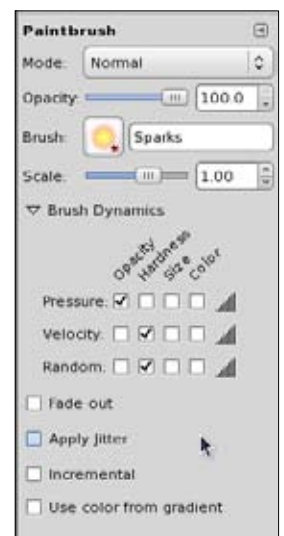


Figure 3: Brush dynamics



Figure 4: New text selection and management tool




Figure 5: Pan beyond border




Figure 6: GEGl operations

yet finished the integration. GIMP, by default, still thrives on 8-bit colour support. There are no documents available for GEGl, or to be precise, the documentation team hasn't yet released documents for the new version.




Figure 7: Different shapes in GIMP




Figure 8: Colour select tool in action

You can activate GEGl from *Colours*→*Use GEGl*.

There are only two operations possible with GEGl—one, to enable it, and the other can be accessed from *Tools*→*GEGl Operations*. This option has lots of options like Gaussian Blur and White Balance, to name a few.

Additionally, GEGl operations and plug-ins are sluggish as compared to old plug-ins. A newbie will surely get confused with what GEGl is all about. Though the GIMP 2.6 launch announcement does promise better inter-portability of GEGl and with version 2.8.

Apart from these major changes, there have been modifications in existing tools. So let's take a look at some of the better plug-ins that come along with GIMP 2.6

GIMP selection tools

One of the unique tools that can give Photoshop a run for its money is the selection tool provided by GIMP. Well, of course, PS also has selection tools. However, what makes GIMP tools unique? The ability to resize even after making a complete selection. In addition to this, GIMP's rectangular selection tool has the ability to add curved edges in a few clicks, which makes your

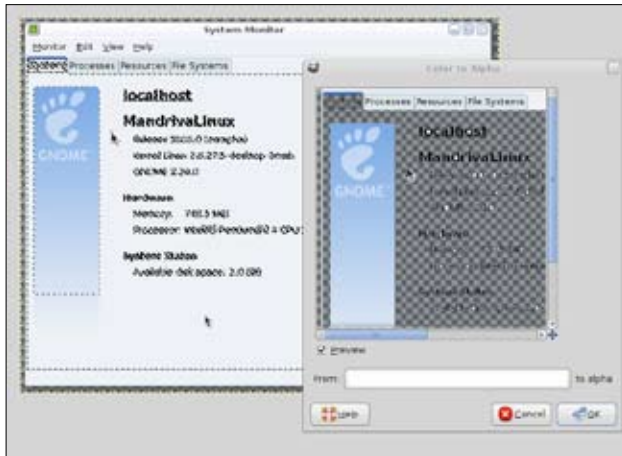


Figure 9: Fuzzy Select and Colour to Alpha

job much snappier and more efficient while creating headers, buttons and logos. The new polygon selection tool is the real icing on the cake.

Colour selection and 'colour to alpha'

If you are not well versed with the Path Tool and Free Selection, and you want to remove the background from an image, then these two colour tools come in handy. Both of them work pretty differently.

With the colour select tool, you first select the tool from the tool bar and then click the colour you want to select/remove, and press *Delete*. That will remove the selected colour from the background. But it will leave some traces on the border of the image/colour. Later on, you can use the blur effect on the borders to smoothen it.

'Colour to alpha' takes the same process to the next level. It makes it even easier for you to remove the background. Select the colour you want to remove from the 'wand' tool and navigate to *Colour→Colour to alpha* and apply the effect.

Custom brush

Have you ever wondered how those beautiful brushes are made? You probably think that one must be adept with vector graphics to create these brushes. Well, you are wrong! Although most of the scalable brushes are created using vector graphics, creating brushes in GIMP is child's play. Yes, simply open an image, select it and copy it, and bingo! You have just created a new brush. This new brush will appear in the brush window and you can save it by clicking on the *Save* button.

The GIMP is certainly one of the best image editors available out there, and version 2.6 is no exception. Though a majority of the changes lie in the core, some of the user-level features make it a pretty nifty update. The inclusion of GEGL and CMYK support has certainly opened new doors for this venerable FOSS image editing application.

Version 2.6 tries to leverage the gap between paid alternatives and freely available tools. But with incomplete

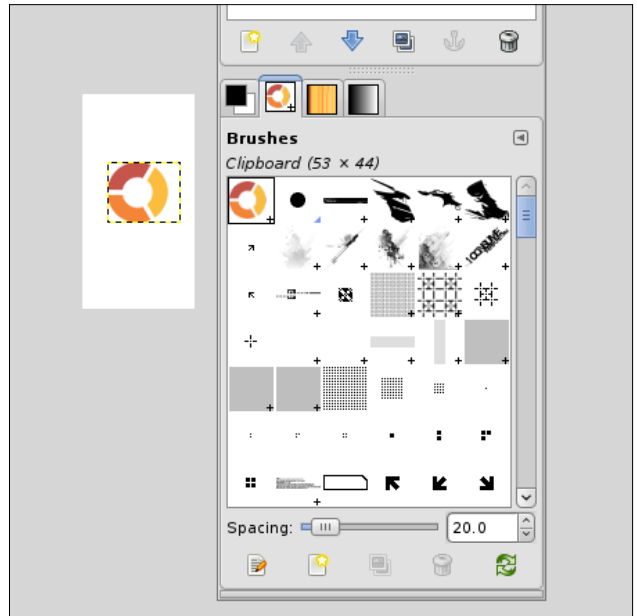


Figure 10: Creating a new brush

integration and UI improvements it still has a long way to go.

Looking forward, it seems that version 2.8 will be a killer release. Developers have revealed some plans for the next stable version: on-canvas text editing, merger of finished projects from the Google Summer of Code, a revamped user interface and better GEGL integration.

Till then, let's wait. But, just in case you haven't tried GIMP out, then 2.6 is the release for you.

As for the missing features... well, as they say, "Rome was not built in a day!" The GIMP has already gained some much-needed momentum, and with GEGL and CMYK on board, it seems it's destined to gain a lot of respect among the pros. **END** 

The GIMP 2.6.3



Pros:

Fast, uses low resources, great alternative to Photoshop, FOSS, GEGL, lots of plug-ins and tools available.

Cons:

Complex user interface for newbies, lacks few advance image editing features, incomplete GEGL integration.

Platform: Universal

Price: Free (as in beer)

Website: www.gimp.org



Notes and references

- Home Page: www.gimp.org
- Release Notes: gimp.org/release-notes/gimp-2.6.html
- Gimp Wiki: en.wikipedia.org/wiki/GIMP
- GIMP UI: gimp-brainstorm.blogspot.com
- CMYK: en.wikipedia.org/wiki/CMYK

By: Shashwat Pant

The author is a FOSS enthusiast interested in Qt programming and technology. He is fond of reviewing the latest FOSS tools and distros.



A Multi-booting Template

With so many Linux distributions around, each with its own flavour, it would be a waste not to sample them from time to time. Here is a simple multi-boot template for your PC for an arbitrary number of distributions controlled by two levels of boot menus.

People often ask why there are so many Linux distributions. No one seems to ask why there are so many items in a buffet spread or so much variety in a botanical garden. The large number of distributions sure makes it difficult to pick the right one for a particular task; but having picked the right

distribution, you have it better adapted for the task at hand.

You may not be able to sample all the items on a buffet, but you certainly can multi-boot your machine with a bunch of lively Linux distributions and zero in on your favourite set or—even better—keep all of them.

Partition scheme						
sda1 Primary	sda2 Primary	sda3 (Extended)				sda4 Primary
		sda5 logical	sda6 logical	sda7 logical	sda8 logical	
/boot	swap	/tmp	/ (distro 1)	/ (distro 2)	/ (distro 3)	/home
100MB	2GB	2GB	12GB	10GB	10GB	84GB

Table 1

Look at it another way. You share a couple of Linux (of course!) PCs at home with family members with diverse requirements: your grandmother wants simplicity, the cousin wants multimedia, the kids want games and education tools, and you want to try out the latest distributions. Now you can meet all these requirements and pamper yourself in the bargain.

Let us get straight into looking at a partition and multi-boot scheme that you can use as-is or tailor to your requirements. I'm assuming that you understand the basics of partitioning, boot loaders, and have installed Linux distributions in the past. Even if you haven't, I am sure you can get a leg-up from the neighbourhood know-all.

Template

Our partition and multi-boot scheme has a two-level boot menu. It has a master boot menu listing all the installed distributions available for booting. Once you select the distro to boot from, you are taken to the dedicated GRUB menu of the specific distro you are booting.

For instance, when I select 'openSUSE 11' on the master boot menu, I am taken to the SUSE boot menu that lists the various options to boot openSUSE such as: a 'normal boot' or a 'fail-safe session'. We will revisit boot menus later in this article. For the moment let us get down to partitioning our disk. We are assuming a blank hard disk to keep things uncomplicated. You can start the partitioning session from a live CD such as Knoppix.

Start with a primary partition (sda1) of about 100 MB for `/boot` followed by a primary partition (sda2) for swap of a suitable size. Make sure the swap is larger than your RAM because this is the space used by the suspend-to-disk feature on some distros.

Of the remaining two primary partitions possible, make one an extended partition. This extended partition (sda3) will house the following logical partitions:

- A partition (sda5) to mount as a common `/tmp` directory for all distributions. About 2 GB for `/tmp` should do for most people, but again this is a suggestion and not a recommendation.
- Slots (sda6, sda7, and so on) for the various distributions. About 10-12 GB per distribution should be adequate for most distributions and you could even do with less.

The fourth primary partition (sda4) can be mounted as the common `/home` directory for all distributions. A common `/home` across distributions can lead to some irritations related to the desktop and we will explore a workaround later when discussing usage scenarios. If you have a capacious hard disk (and who doesn't these days), you can have more than one partition available for mounting as `/home` and specify which one to mount automatically for a particular distribution during install time. If you have more than one partition for `/home`, make sure you create the additional ones as logical

Partitions and mount points						
Distro	/boot	/tmp	Swap	/	/home	GRUB
1st Distro	sda1(F)	sda5(F)	sda2(F)	sda6(F)	sda4(F)	In MBR
2nd Distro		sda5	sda2	sda7(F)	sda4	In sda7

Table 2

partitions inside the extended partition sda3. Choosing a common swap and `/tmp` partition across distributions helps you utilise your hard disk real estate that much better. You have a choice of filesystem types to use for each partition, but remember that GRUB is finicky and prefers ext3.

Depending on how much space you need for swap, `/tmp` and `/home`, you can easily arrive at how many distributions you can pack into a single hard disk. As a variation you can also pack all distributions on one hard disk and keep your data on another.

Table 1 shows how the partitions might look like for a single disk of 120 GB with three distributions sharing a common home partition.

Install procedure

One approach is to start with a bootable GRUB disk and install GRUB in the MBR of sda. This way the boot partition is not associated with any distribution. But let us take a simpler alternative and let the first distro install the boot loader for us. This way, the dedicated `/boot` partition gets associated with the first distribution but that does not make any material difference to the boot interface.

To keep things systematic and under control, start by installing the first distribution. Using Table 2 as a cue, install the first distro by providing the indicated mount points for the respective partitions. For this distro you will install GRUB in the MBR and use sda1 as `/boot`. For all other distros, GRUB goes into the "first sector of the boot partition". The (F) next to each partition indicates it needs to be formatted. Note carefully that from the second distro onwards, you need to format only the `/` partition. Remember this, lest you end up losing data that already exists in the `/home` or `/tmp` partitions.

Configuring the GRUB menu

Since we are letting the first distro take charge of the boot partition and install GRUB for us, all we have to do is edit the GRUB menu to add the other distros. The GRUB menu exists as `/boot/grub/menu.lst`. Here is what mine looks like today for four distros:

```
# grub.conf generated by anaconda
#
# Note that you do not have to rerun grub after making changes to this file
# NOTICE: You have a /boot partition. This means that
#         all kernel and initrd paths are relative to /boot/, eg.
#         root (hd0,0)
#         kernel /vmlinuz-version ro root=/dev/sda6
```

```
# initrd /initrd-version.img
#boot=/dev/sda

default=1
timeout=10
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu

title Fedora 9 (2.6.25-14.fc9.i686)
    root (hd0,0)
    kernel /vmlinuz-2.6.25-14.fc9.i686 ro root=UUID=xxxxxxx-xxxx-xxxx-
xxxx-xxxxxxxxxxxx rhgb quiet
    initrd /initrd-2.6.25-14.fc9.i686.img

title OpenSUSE 11 @ /dev/sda7
    rootnoverify (hd0,6)
    chainloader +1

title Mandriva 2009 @ /dev/sda8
    rootnoverify (hd0,7)
    chainloader +1

title Linux Mint 5 (Elyssa) @ /dev/sda9
    rootnoverify (hd0,8)
    chainloader +1
```

Briefly, the initial file was generated by Fedora 9 when I installed it as the first distro. Adding the subsequent entries for other distros is easy. We just chain-load and transfer control to the GRUB of the distro we want to boot. Each of the distros, in this case openSUSE, Mandriva and Mint, automatically create their own attractive GRUB menus at install time and we can leave them untouched. Remember to get the boot partitions right for each GRUB menu entry.

Sharing the home directory

To share your files seamlessly across your Linux installations, create your users in such a way that they have the same UID on each distribution -- each with their own dedicated group. For instance, you could have the uid:gid series of 500:500, 501:501, and so on. This is the simplest scheme and should work well for a typical home-computing environment. You can also create users in such a way that certain users have access to certain distros and no access to others. You can thus restrict each user to just one distribution if you wish.

Typical usage scenarios

The default scheme with the common home directory across distros works fine, except for some side effects on the desktop. Since desktop (KDE, GNOME, Xfce,...) versions vary across distros, you might get widget load errors and some unintended desktop icons while hopping from one distro to another. This is because the 'dot' files for your desktop are common across all distros. There is

also the possibility of opening a file with a newer version of a software in one distro and rendering it unreadable by the same software with a different version number from another distro. As a solution, the following two approaches have worked for me, singly or in combination.


Rather than mount a common */home*, you can let the home directory stay with the respective distro so that your desktop stays there as well. Each distro then gets its own desktop and there are no unpleasant side effects. You can then mount your previously designated home partition separately under, say, */mnt/home*. If you like this approach, provision adequate space for the root ("/) partition of each distro (so that a 'small' */home* directory can stay there) and do not give any partition as a mount point for */home* while installing. This way the */home* directory gets created in the root partition of the distro, automatically.

The second approach is to follow some discipline. It helps to have a main distribution that you will use for mail review and other daily chores. You would rather not be messing around with mailboxes across distros. Having voted for your main distribution, you can then boot the distribution that specialises in the task at hand: perhaps media editing or gaming. With this approach, specific file and data types stay associated with a specific distribution and sanity can prevail all round.

Your mileage will vary depending on the distros you multi-boot. You should experiment to arrive at the template that works best for you and the work discipline that goes with it.

Pay day

Going by the examples in the references, you will not reach the limit on the number of distros that can be booted with a similar scheme any time soon. Once you have your PC set up with a flexible partition scheme, evaluating and comparing new distros becomes a walk in the park.

The grapevine has it that this leads to a condition called distromania. Symptoms include staring for long hours at installation progress bars and then rushing to boot a new distribution every couple of hours. Some people may need to be physically pulled away from their computers. Fortunately, in my case the condition was self-limiting—after trying out all distros that came with *LFY* in the past few months, I realised I had run out of CD blanks to burn more ISO images of lesser known distros. Then I quickly got down to writing this article before the next affliction could hit me.  **END**

References

- www.justlinux.com/forum/showthread.php?t=144294
- www.justlinux.com/forum/showthread.php?t=147959

By: Gurudutt Talgery

The author is a FOSS enthusiast and practitioner who never stops marvelling at the liberating ideology of FOSS.



Improve Multi-OS Computer Performance ...through Cross Swapping

Computers with more than one hard disk and more than one operating system (OS) can be set up to provide better performance than what OS installations generally default to. This technique, called cross swapping by the author, is described using a dual-boot scenario of Linux and Microsoft Windows.

*L*et's first understand what virtual memory is. Today, computers are used for a wide variety of tasks, facilitated in part by modern OSs providing multitasking capabilities. A multitasking OS allows you to run more than one program at a time. If the number of programs you have started is greater than the number of CPUs (processors) available and in use, the OS repeatedly runs tiny parts of each program for very small amounts of time (not noticeable by the user) by turn; that is, it repeatedly allows each program to use the available CPU(s) for tiny fractions of a second. This creates the illusion of all programs running in parallel.

Multitasking, however, places large demands on the system memory (RAM). Running several graphical programs

simultaneously can easily exhaust even today's large memory banks in a short while. To alleviate this problem, modern OSs move some unused program data from RAM onto the hard disk to make room for required data.

This can be done because all of a program's data in RAM is not needed at all times. If a program has finished working with some data and needs other data that had earlier been moved out to the hard disk, the OS can again move out the used data and move the old data back in. This entire process of moving unneeded data out and required data in is called swapping or paging, and the area on the hard disk where data is moved out from the RAM is called the swap space, page file, or virtual memory. Virtual memory is used even if you do not run more than one program at a time, simply because

it allows the OS to keep frequently used parts of the running program in memory (as well as parts of the OS itself).

It is always better to assign a dedicated partition on a hard disk to the virtual memory. Since hard disks are much slower than RAM, it takes a significant amount of time to move data to and from virtual memory. Having a separate partition causes reading and writing of virtual memory to be independent of the reading and writing of other files. We will see how this scheme can be improved further.

Improving performance through cross swapping

If you have more than one hard disk, you can improve the performance of the virtual memory further. To do so, simply install the OS on one hard disk and the virtual memory on a dedicated partition on another hard disk. This scheme is much better than having both on the same hard disk because it gives independent control to the virtual memory. While files are being read from or written to one hard disk, the OS can operate the virtual memory simultaneously on the other—it won't have to wait for either operation to complete before starting the other.

These days, it is becoming increasingly common to have more than one OS installed on a computer and use (boot) any one of them at a time. If you have more than one hard disk, you can use cross swapping to utilise all hard disks optimally.

It is best to install separate OSs on separate hard disks. This isolates their partition tables, provides physical encapsulation, and prevents space conflict and re-installation if you want to change partition sizes later. However, typically, only one OS is used at a time. All the others remain unused as long as that one is being used. If the other systems are on different hard disks, those hard disks will remain unused for that time. We can therefore utilise them for the virtual memory of the OS that is running. To do this, simply allot each OS's virtual memory to a partition on a hard disk other than that on which the OS is installed. We will look at a scenario where Microsoft Windows and Linux are to be dual-booted.

Let's suppose we have two hard disks, A and B. We install Windows on hard disk A and a distribution of Linux on hard disk B. Now, to employ cross swapping, we place the virtual memory for Windows on a dedicated partition on hard disk B, while we place the virtual memory for Linux on a dedicated partition on hard disk A. However, to achieve this, it is best to create the partitions before or while installing the operating systems. They can be created later, but that requires either a RAID/LVM set-up or special software. If you wish to follow the process mentioned, set up your hard disk partitions as shown in Figures 1 and 2.

Figure 1 shows hard disk A with the first two partitions allotted for Windows files and the last partition used as dedicated virtual memory for Linux (the swap partition). Figure 2 shows a similar arrangement of hard disk B with a virtual memory partition for Windows and the last two partitions for Linux files. The exact positioning of the

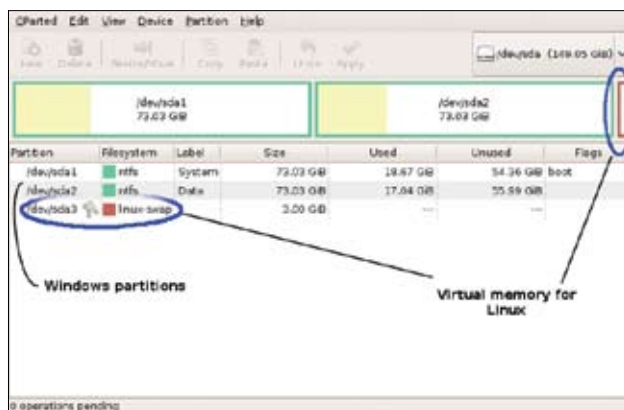


Figure 1: Hard disk A with Windows partitions and a virtual memory partition for Linux

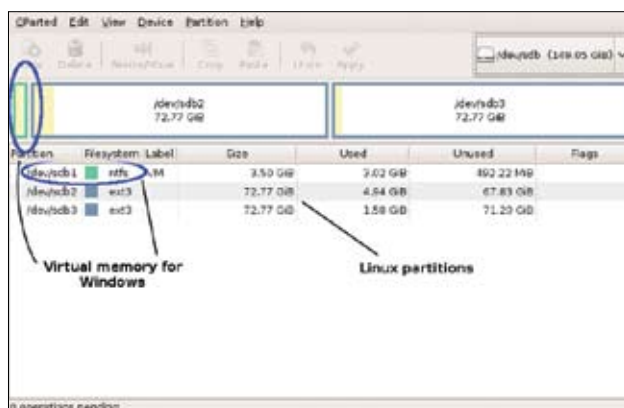


Figure 2: Hard disk B with Linux partitions and a virtual memory partition for Windows

partitions is such because Windows always needs the first partition of the first hard disk, and the Windows swap partition is placed at the head of the second disk because Windows sometimes achieves better performance if its virtual memory partition is placed at the beginning of a hard disk. The two partitions each for Windows and Linux files are just my personal preferences to keep the system and personal files separate. You can use one or more partitions in their place for each OS.

To actually create such a layout during installation, start with two blank hard disks and follow the steps below (note: this assumes we're using Windows XP and so the exact steps may differ for other versions—refer to Windows documentation for changing the amount of virtual memory in such cases).

1. Create a primary partition on the first hard disk, make it bootable (active), and install Windows on it (you can create the partition from within the Windows installation program itself).
2. If you want another Windows partition for your personal files, as shown in Figure 1, create a primary partition on the first hard disk (Disk 0) using the Windows disk management tool (right click on *My Computer*, click *Manage*→*Disk Management*).
3. Leave the area for the Linux swap unpartitioned at the end of the first hard disk (Disk 0). Keep about 1.5 times the amount of your RAM reserved for this space. You

can also keep 1 GB if you have 512 MB or more of RAM and want to conserve disk space.

4. If the disk management tool is not already open, right click *My Computer*, click *Manage*→*Disk Management*.
5. On the second hard disk (Disk 1), create a primary partition 1.5 times the size of your RAM (or 1 GB—see the above point); plus about 15 to 20 per cent of that space, assign a drive letter and format it (preferably with NTFS). Note that only Windows needs the extra 15-20 per cent of free disk space. Don't reserve any extra amount for Linux on the first disk. Close the disk management tool when done.
6. Right click *My Computer*, click *Properties*→*Advanced*. Under *Performance* click *Settings*→*Advanced*. Under *Virtual memory* click *Change*.
7. Select the first partition (C:), click *No paging file*, then click *Set*.
8. Select the new partition (drive letter) you created in Step 5, then click *Custom size*.
9. If you allocated 1.5 times the amount of your RAM, type that amount in megabytes (1 GB = 1024 MB) in both the minimum and maximum boxes, then click *Set*. If you allocated 1 GB, type 1024 in both boxes, then click *Set*. This sets up your virtual memory for Windows. Make sure not to set the whole partition space (which includes the additional 15-20 per cent space) here. The extra free space is required by Windows to perform optimally.
10. Click *OK*, then *OK* again to close the *Performance Options* box.
11. Click the *System Restore* tab, then select the partition (drive letter) you set the virtual memory in.
12. Click *Settings...* then select 'Turn off System Restore on this drive' and click *OK*. Confirm that you want to turn off *System Restore* on the drive.
13. Click *OK* to close the 'System Properties' box.
14. In your Linux installation program, select the option for custom (or expert) partitioning at the beginning of the partitioning phase (most Linux installers provide such an option, but if yours doesn't, you can use the *fdisk* or *cfdisk* program; (*cfdisk* coverage is out of the scope of this article so you can get help from their manual and info pages on a working/live Linux system).
15. Select the first hard disk and create a primary partition in the space you had reserved for Linux swap at the end of the first disk. Make the partition use the full reserved (available) free space. Set it as the swap partition. This creates the virtual memory for Linux.
16. Select the second hard disk and create primary or extended partitions as you want for your Linux files after the existing Windows partition. The existing Windows partition on the second disk contains the Windows virtual memory that was set up previously.
17. Proceed with your Linux installation. Don't set the option to automatically mount the Windows virtual memory partition on the second hard disk. If your installer turns it on by itself, make sure to turn it off.

Some notes

The exact performance gains you enjoy will vary depending on the amount of RAM you have (the less the RAM, the more the virtual memory used), the number of applications you use at a time and the load they place on the system memory, and other similar factors. However, even if not perceptible, there will be some difference, which will be better than using the default virtual memory set-up. You can check the amount of virtual memory being used, by using the *Task Manager* in Windows and the *System Monitor* in Linux.


A maximum of four primary partitions is allowed on a hard disk. More than four partitions can be created by using logical partitions within an extended partition. If you want to use logical partitions for your Windows files, fill the first hard disk to the end with the extended partition, create your logical partitions and make sure to keep the Linux swap space reserved within the extended partition. Then, in Step 15, create a logical partition for the swap space instead of a primary partition. Note, however, that this is not the optimal set-up for virtual memory.

If you have more than two hard disks and want to install more than two operating systems on them, you can jumble up the virtual memory partitions for them in a similar manner. However, note that cross swapping can be implemented best if you have as many operating systems as hard disks and install each OS on a different hard disk. Also ensure that you keep the first partition on the first hard disk for Windows.

If you use IDE hard disks, the primary master is the first disk, the primary slave the second, the secondary master the third, and the secondary slave the fourth. However, for best performance, you should make your hard disks masters on separate IDE channels (e.g., primary master and secondary master). For SATA hard disks, the disks are ordered according to the SATA ports on which they are connected.

Windows does not recognise/read Linux and UNIX partitions; so it is safe to place the swap partition for Linux/UNIX on the hard disk on which Windows is installed. However, do not try to manipulate any Linux/UNIX swap or other partitions using the Windows disk management tool.

You can allot partitions for the virtual memory of an OS on more than one hard disk with one partition on each hard disk. This will improve performance further. However, don't allot a virtual memory partition for an OS on the hard disk on which the OS itself is installed, else you will defeat the purpose of cross swapping.

If you use more than one Linux distribution, you can reuse the same virtual memory partition(s) for all of them. 

By: Saurav Sengupta

The author is a final year IT student and a software developer with an interest in both application and system programming. He can be reached at sauravsengupta01@gmail.com.

Enabling Indian languages on the FOSS desktop Part 3

Your Desktop, Your Language

In the past, we've looked at how to enable word processing in your native language. In fact, your entire desktop, menus included, can be displayed in your language. This process of translation is called localisation. Find out more about it.

From operating system commands to menu entries, almost every new computing term is in the English language. The age-old problem for translators is to convey the collocations, and *moods*—baggage that doesn't often translate too well. Software translators, therefore, must be particularly careful in dealing with such terms, for fear of producing awkward and stilted work, which could mar the FOSS experience. This is the paradigm within which we'll critique the Indian-language FOSS desktop and gauge its state of readiness for you, the user.

Naturally, I can only comment on the languages that I know, and I pick Hindi for this review. I installed support for the language in my Ubuntu 8.10 GNOME desktop. *Kaisa tha*, then?

Installing and activating local-language menus

Very few distributions have Indian language menus installed by default. In Ubuntu 8.10, install the feature yourself, as follows: *System menu*→*Administration*→*Language Support* (Figure 1). Now, in the window that appears, select your language, and Ubuntu will download the fonts and menu files. It's a large download! Make sure you have a fast Internet connection. This step will also install your language keymap, so that you can word-process in your language. So if you can type in your language as the earlier article in this series described, then you've already got translated menus installed. Neat, huh?

But you'll have to activate the menus. Log out. Now, in the login window, select the *Options* menu and then *Language*. Choose your language here. Log in again with your user name and password. Hey

presto, *aapke menus ab Hindi mein bhi!* (Refer to Figure2.) *Kya khoob baat hai...*

Khidki or window?

...or is something else? Most of the common e-terms such as *open*, *save*, *view*, *edit* and *rename* are translated competently. One place where the translators have excelled in particular is '*Raddi mein bheje*' (of course, it's written in Devanagari script in the menus!) which is a delightful and comprehensible translation of 'Send to Trash'. It translates the concept of a trash can into that of a pile of *raddi* on the image level. Way to go! Immigrant concept, naturalised fully.

But a very few words, though not quite lost at the airport, are delayed by the customs. Consider 'page layout', and then ponder over the Hindi translation in the OpenOffice.org suite of Ubuntu 8.10—*prushtha vikhandan purvavalokan*. I'm no Hindi expert, and presumably more people than Ashutosh Rana speak like that—no offence intended—but everyone doesn't. And I have never *adhithapit* (install) any *vancchit* (default) applications. These concepts had better be trans-created instead of translated. They look really awkward in the guest language.

These are rare, very rare (and comic) issues with the Hindi desktop. The translators are not entirely to blame. What can they do, when for example, there is no adequately co-locational term for 'window' in Hindi? The word 'khidki' cannot be suggested except in humour, until the word accumulates new meaning. Dear, dear me. A bold alternative could be to translate the concept of a 'window' using an Indian visual (as was done while translating the 'Trash Can' concept). '*Jharokha*' here could be a choice. It has co-locational pedigree as well as a really nice sound.

The desktop menus, right-click menus, dialog boxes and mouse-over info have been translated, as well as the menus of the popular applications such as OpenOffice.org writer, spreadsheet, presentation, Evolution mail, dictionary, movie player, Pidgin instant messenger, Gedit, Terminal, and GNOME games. Mozilla Firefox is not touched. As for the Rhythmbox music player, it is only partially translated. Information and warning messages are also translated. And files can be renamed in your language script also. Rather bafflingly, the help files have not been translated at all!

At its current state, the desktop in our test distribution looks as if a native speaker would be at home with it. Oh, how he can *vikhanditise* his *prushts...*

Translate your own desktop

Eleven Indian languages are being worked into KDE4; GNOME has 15. There's lots of work still to be done. For example, the FOSS Hindi desktop in GNOME 2.24 is 80 per cent translated. The KDE4 desktop is 70 per cent translated, according to their websites. And maybe there are translation issues in a couple of places. For example, very '*propahi*' Sanskritised Hindi is used, but that is not the only Hindi spoken in India. There's scope for translators



Figure 1: List of supported languages





Figure 2: My localised desktop—not all items are translated to Hindi as you can see

who want to see 'their' language represented on the desktop. That must be the case for many Indian languages.

So, how do you get involved? You link up with other volunteers, that's how. Visit the KDE and GNOME websites; they have separate sections for localisation. Or you can visit your distribution's website and join their team. For Ubuntu, visit the Ubuntu-India website at ubuntu-in.info. For Fedora, your place is translate.fedoraproject.org. And so on, for your distribution.

Generally, the process of localisation is this: the material to be translated is uploaded on these or related websites, and you have to translate it in a special editor. That done, you submit your work by e-mail or upload. Simple! Not at all time-consuming either. It's a great way for non-techies to contribute to FOSS, isn't it?

Sign-off

Localisation is associated with a fair bit of hype; but the FOSS localisation communities, especially the Indian ones, if Hindi is any indication, are actually implementing it in a usable way. 

By: Suhit Kelkar

The author is a freelance journalist and translator based in Mumbai, India. He can be contacted on suhitkelkar@gmail.com

sabayon⁴

Is It Ready For You?

It's arguably the most feature- and multimedia-rich distro out of the box...and it's now available in a new avatar. With massive improvements in the software database and in hardware support, the much-hyped Sabayon version 4oh is knocking on your doors.

*S*abayon, formerly known as RR4/RR64 Linux, has been built on the unstable branch of Gentoo. Sabayon 4 is the latest offering from the Sabayon Terminal. 4oh! is a rolling release, which means you don't have to reinstall the OS if a new version is released; easy updates will be made available for renewal of the OS.

Soon after the initial release, developers released the first revision for Sabayon—i.e., 4r1. This first revision, now bundled with this month's magazine, has a lot of bug fixes and adds a few newer versions of the available software.

Sabayon 4 has not only introduced a rejuvenated operating system, but also a website that makes sure that you won't get lost in finding what you're looking for. The new interface is eye-catching, less cluttered and gels well with the Sabayon 4 artwork.

Sabayon 4 is a major leap from version 3, as you will observe significant improvements over the previous stable version. After two years of exhaustive planning and development, developers finally came up

with a novel version to enchant audiences with something new and potent. Although it looks to avail all the latest technology, will the new version measure up to its name? Will it be able to uproot the big players? That's exactly what we're here to find out.

My test system had the following specs:

- AMD Phenom X3 8650 overclocked to 2.7 GHz
- 4 GB DDR2 800 MHz RAM @ 1000 MHz
- 640 GB Western Digital Aaks
- AMD 790GX with ATI HD 3300 graphics
- HP's all-in-one 1410 printer
- A 32-bit edition of version 4.0 r1 Oh!

Get, set, go!

After popping in the live DVD you will be greeted with the sheer simplicity of the menus and the impressive boot screen theme. The DVD menu comes with lots of options for booting and installation. You can use it as a live DVD, directly install it on the hard drive and much more... Apart from this, Sabayon allows you to play high-end games without even leaving the live mode. Booting the OS is quite smooth. During the

boot-up you will note something quite unique. Sabayon developers have added music that plays while booting into the live mode only. So, even if you have a low-end PC, you still won't be bored by the old boot screen's progress bar. After the bootsplash you will be asked to select the desktop type based on desktop effects. It not only allows you to switch between Compiz and a simple desktop, but also notifies you about your graphics adapter's compatibility with Compiz.

Finally, the desktop showed up. Figure 1 shows what the default desktop looks like. My initial impression was one of mild disappointment. Except the artwork, everything else was nearly the same as that of the 3.x edition. The placement and the number of icons haven't changed yet. The default desktop looks haphazard—having so many icons, I believe, can confuse a newbie.

As far as its hardware detection capabilities are concerned, it was unable to configure my resolution correctly. It chose to set it to a resolution of 1024 x 768 pixels instead of 1440 x 900.

I took a sneak peak at the installation options quickly, but there was nothing new to report. Sabayon uses the good-old Anaconda as its installer. It has added nearly all the major desktop environments (Figure 2), and has a specific installer option for Netbooks, which looks like a handy addition if you are using such hardware. The installer automatically adds KDE if you select GNOME, and vice versa.

The installation process is pretty awkward. Initially, it even installed KDE, although I had deselected it. However, at the end of the installation, it removed the deselected packages. Quite a waste of time, I'd say. The installation of deselected packages could have been skipped. In my system it took a total of 23 minutes. I personally have no issues with the time, as Sabayon installs quite a lot of packages. If you select all the options, then you might end up installing more than 12 GB of software.

Into the drive and what do I find?

Even after installation, Sabayon didn't set my resolution correctly despite the presence of ATI Catalyst drivers out of the box. I tried fiddling around with *xorg.conf*—even adding a new depth section and new resolutions, but to no avail. I was shocked that I couldn't use the *aticonfig* command even after the drivers were installed. Whatever I try, I just couldn't set my resolution correctly. Surprisingly, Mandriva using the same proprietary ATI driver (*fglrx*), sets everything correctly. Although it was unable to set my resolution correctly, it was able to set my graphics card to run Compiz right away (Figure 3).

Since Sabayon comes loaded with all the printer drivers and utilities, I thought it's time I set up my printer. I own a HP all-in-one 1410 printer, which has pretty good Linux support. I have used this printer with all the leading distros without any problems. To use it in Sabayon, I first opened the HP Device Manager. Although



Figure 1: Sabayon's default KDE 3.5.x desktop



Figure 2: Desktop selection screen on Sabayon's hard disk installer



Figure 3: The 3D desktop out-of-the-box

this application did detect my printer, it failed to add the respective drivers. So I manually selected the driver from the vast HP database and, to my surprise, in following the series of 'Next' steps, I got a prompt saying, "Printer not configured". This really irritated me, and I launched



Figure 4: Sabayon's GNOME desktop

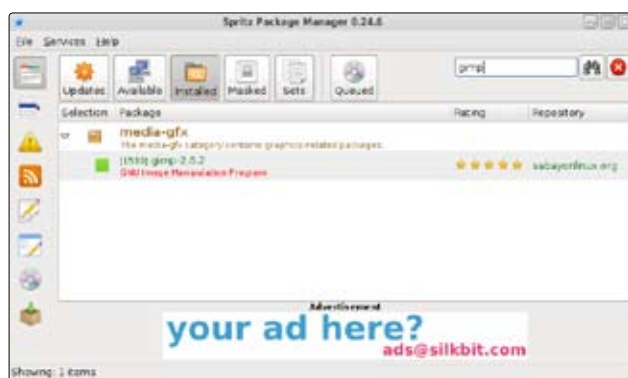


Figure 5: Spritz package manager

the GNOME Printer Manager. This application was able to configure my printer correctly for a change. However, when I tried to use XSane to test-scan something, it again reported that there was no device found.

I didn't expect Sabayon to have such issues with my hardware, since it included the latest OSS tools. Other distros that utilise the same tools detect and configure these easily.

Considering that Sabayon comes with kernel 2.6.27, which includes drivers for Atheros Wi-Fi chips that supposedly provide advanced support for the latest Wi-Fi routers, I think there shouldn't be any problem using it. I didn't have one. See if you have any issues and do report your experience.

Ah, all these goodies...

Application availability is an area in which Sabayon really shines. In fact, when it comes to out-of-the-box competition with other distros, Sabayon wins hands down (apart from Knoppix, perhaps). The DVD is 4.13 GB, which pretty much expands to a requirement of more than 12 GB for a full install. So do make sure you have enough space before a full installation. Whether you need an application for your day-to-day activities or some really exquisite feature, it's all pre-loaded! A few of the basic software and core utilities included in the OS are:

List of applications	
Software category	Available applications
Accessories	Archive Manager, LXTerminal, Xarchiver, Kworldclock
Games	Battle for Wesnoth, Nexuiz, Sauerbraten, Second Life, KDE and GNOME Games Pack
Graphics	DJView, FortForge, F-Spot, The GIMP, Gpic View, Hugin Panorama Creator, Inkscape, OpenOffice.org Draw, X-Scane, Kooka, Kfax
Internet	aLink Creator, aMSN, aMule, Deluge Torrent Client, Ekiga VoIP, Epiphany Browser, Filezilla FTP, Google Earth, Konversation, Logjam, Firefox, NXClient, Pidgin, Seamonkey, Skype VoIP, Wxcas, Zenmap
Office and productivity	Evolution, OpenOffice.org Office Suite
Programming	Beanshell Prompt, Cmake, Doxy Wizard, Glade Interface Designer, iLConstars, Monodoc, Qt Designer, Quanta Plus, KHTML Image Map Editor, Klink Status, KXSLDgb
Sound & multimedia	Ardour2, Audacious, Audacity, DVDRip, Elisa Media Centre, K3b, K9Copy, Lacie Lightscribe Labelers, Lightscribe Simple Labelers, Totem, Mplayer, SMplayer, TV Time, VLC, Xine
Education	KDE education package
Systems tools	ATI Control Centre, Avahi Zeroconf browser, Bluetooth File Sharing, CUPS Utils, HP Device Manager, Nvidia Control Centre, Virtual Box, VirtManager, Wine-Doors, Wireshark

Table 1

1. KDE 3.5.10
2. GNOME 2.24.2
3. The GIMP 2.6.2
4. K3b 1.0.5
5. Ekiga 2.0.12
6. OpenOffice.org 3.0.0
7. Firefox 3.0.4
8. Skype

Refer to Table 1 for a list of some of the available applications.

The application versions are pretty much up-to-date. However, one thing to note here is that it includes a poorly-configured GIMP. The present version mimics the window management behaviour of the older version 2.4, and even editing the settings from *Preference-->Windows Management* didn't restore the original window behaviour.

KDE4 didn't make it into this release, but you can install it from the online repositories, and that was the main reason for me to install GNOME over KDE 3.5. Although, the KDE4 integration is at par with KDE 3.5, if you are a GNOME user you will be a bit disappointed with the way it's been integrated. The menus are cluttered and applications are scattered across wrong sections. You will find graphics programs under music and vice versa. Figure 4 gives you a glimpse of the default GNOME desktop offered by Sabayon.

It is always good to have alternatives, but how many alternatives will you use at a time? Sabayon includes four Web browsers, nine media players, multiple archive managers, image viewers and other tools. Having too many similar tools makes the system bloated and the OS memory voracious. The distro should have made some shrewd choices on packages. Fewer and better packages might have resulted in a snappier experience.

Sabayon 4 has added a few great utilities. It includes full-blown virtualisation solutions from the geeky Virt-Manager to the user-friendly Virtualbox. Sabayon also includes the Lightscribe media labeller. This tool allows you to create covers for your lightscribe CD/DVD in just a few clicks. You can then use the software to print it on the discs using the Lightscribe DVD/CD Writer.


As for package managers, it has introduced Entropy, which is fully compatible with Gentoo's Emerge package manager. It has also introduced a graphical wrapper called Spritz (see Figure 5) to cope with other high-end package managers available. Spritz aims to be simple, yet powerful. It is still at a nascent stage of development, so it wasn't a surprise when it crashed quite a few times while testing.

With over 8,000 packages you will find nearly all you desire from Spritz, and can install it easily. Spritz looks and performs nearly

similar to Packagekit, but is still not as powerful. Though I'd say its inclusion is a very positive step and is way simpler than the labyrinthine portato package manager.

Closing up!

Sabayon indeed is a great distro, but version 4 didn't measure up to expectations. The poor hardware support, at least in my case, has left a negative impression. I only hope the upcoming revisions are a lot better and more usable than the current one. However, this doesn't mean that Sabayon won't work for you. Considering the impressive amount of work done to improve on the previous versions, Sabayon is a must try even if you don't want to replace your current distro with it.

Overall, the new version left a neutral impression on me. Will I use it? No, because I can't get certain things to work even after editing the core files. But it may give you some good company, so don't hesitate to take it for a spin, at least.  **END**

Sabayon Linux 4r1



Pros:

Plethora of software, pre-installed 3D games, multimedia support out of the box

Cons:

Poor hardware support, poorly managed menus and applications, complex default desktop

Platform: x86, x86-64

Price: Free (as in beer)

Website: www.sabayonlinux.org

Resources

- Home Page: www.sabayonlinux.org
- Sabayon Wiki: wiki.sabayonlinux.org
- Forums: forum.sabayonlinux.org

By: Shashwat Pant

The author is a FOSS enthusiast interested in Qt programming and technology. He is fond of checking out the latest OSS distros and tools.

NACE
move to the next level

TRAINING & CONSULTANCY

Novell.

GOLD TRAINING PARTNER

BECOME AN INDUSTRY
READY OPEN SOURCE
PROFESSIONAL

LINUX

APACHE

MySQL

PHP

Linux Basics

Linux Administration

Linux Networking

Linux Security

OpenOffice.org 3.0

We Also Undertake Linux Server
Maintenance & Integration

CHENNAI

T.NAGAR : 99405 42222
ANNA NAGAR : 97909 25554
ADYAR : 97909 25550
MYLAPORE : 9003033233
AMBATTUR : 9445322869
PERAMBUR : 93808 76090
TAMBARAM : 9283366552
VILLIVAKKAM : 9940233360
VELLORE : 9994241383
SALEM : 9944175120

PONDICHERRY : 97918 53774

BANGALORE : +91 9940022277

HYDERABAD

MALAKPET : 9849462617
ASHOK NAGAR : 9703333839

www.naceedu.com

Son... this is KDE 4.2 on openSUSE 11.1! Savvy?

...as Captain Jack Sparrow would have said. Well anyway, it doesn't matter whether you want KDE 4.2 aboard or not, but the new gecko can surely set sail your Black Pearl... oops! I mean, your computer.

This time around I tried my hands on the openSUSE 11.1 64-bit version as soon as it was released last December. After a month or so, I'm still using it as the default OS on my PC. So, what caused the switch from Mandriva? Well, it's the openSUSE Build Service [en.opensuse.org/Build_Service]. Build Service, as you know, is this pretty cool service, where newer software packages are made available for a particular openSUSE version—so, you no longer have to wait for

the new OS version to come out in order to check an updated version of an application. My decision was purely based on the fact that KDE 4.2 is around the corner (scheduled for release on January 27), so I didn't want to wait till April for Mandriva 2009 Spring to try it.

In fact, I'm writing this article on KDE 4.2 RC (which was released on January 13) after 'fusing' it with the final openSUSE 11.1 DVD that was released in December. However, before we get into all that, let's see what the vanilla openSUSE has in its bag.



The basics (or plain boring details)

Linux 2.6.27.7, glibc 2.9, KDE 4.1.3 (3.5.10 is also available as an option), GNOME 2.24.1, OpenOffice.org 3.0, GIMP 2.6.2, and more or less all the software updated to their current versions is what the v11.1 DVD is all about. What sets the openSUSE DVD apart from other distros is that the 3.5 branch of KDE is also available for those who still don't want to switch to KDE4 for whatever reason.

The first thing you'll notice during the installation is the brand new 'free' licence, which finally replaces the old Novell EULA. The distro, which is an aggregate of various free software, has been relicensed under GPL version 2. This means the DVD now doesn't contain non-free software like Adobe Acrobat Reader, Flash, and the like, any more. This also means you are not required to select an 'Accept' check box any more, but simply move on after reading the licence (if you so choose).

The rest of the installation routine has remained very similar to openSUSE 11.0 (if not exactly the same—you can read my article on v11.0 in *LINUX For You's* August 2008 issue if you want to know the steps). The other thing to note is the redesigned partitioner. Although I'd hoped that any change introduced should make things a bit easier (or more appealing?), I felt that's certainly not the case here. In fact, the redesign is sort of awkward and may leave you a bit puzzled.

After installation, the boot-up time and overall performance of the desktop is pretty decent, and comparable to the previous version. Here again, there's nothing much to report as we've already talked enough about the recent KDE and GNOME versions, and have seen how they look and feel in the distros that have been bundled with *LFY* since November. However, if I may add, the all-around green openSUSE theme that was introduced in version 10.3 and continues till date can get on your nerves, and I personally would appreciate it if they switched back to the pleasant bluish look and feel.

One thing I'd still like to report is that openSUSE's default selection of applications and the overall integration of the desktop is still top-notch. Although I'd have preferred it if the SUSE folks did something about the Slab menu that they introduced in GNOME with SUSE Linux Enterprise 10 in 2006. Not that I'm saying it's not pretty compared to the default GNOME menu, but when it comes to usability, it's a hideous show-stopper, expecting the user to wait and click several times to find and launch an application. I guess if the GNOME users are looking for a better menu alternative, Gimme [www.beatniksoftware.com/gimmie/Main_Page] is a much better option, with some good features, although it'd still need some love from the UI design team to make it prettier.

Something else that's new with v11.1 is the complete integration of PackageKit (the GUI package manager front-end introduced by the Fedora camp) across desktop environments. Likewise, along with the YaST's software manager, you can also use this front-end (available in the menu by the name of Add/Remove Software) instead. Personally, I'd have preferred if there was only one front-

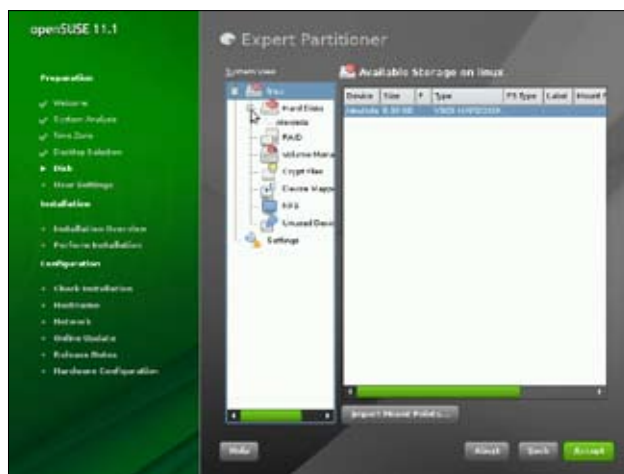


Figure 1: The newly redesigned partitioner




Figure 2: The default GNOME desktop



Figure 3: 3D Alt + Tab in KDE 4.1.3

end available by default. And considering that the YaST package manager GUI for GNOME is so hideous-looking, it'd be neat if the SUSE team focuses on the development of PackageKit full time (to make the GUI better and the software work faster). Not that I'm very comfortable with PackageKit either, as I use the traditional YaST package manager that is accessible from KDE.

Something very handy about the distro is that the package manager automatically installs things like Flash or Fluendo



Figure 4: The default KDE 4.2 RC

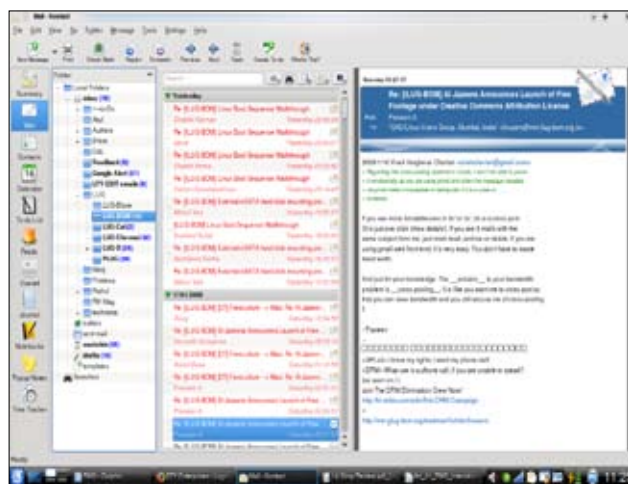


Figure 5: The 'Fancy' theme of Kmail

plug-ins for MP3 support and nvidia/ATI graphics drivers the moment you try and install anything using YaST for the first time. So, if you are a free software purist, you better disable the non-free repository first.

Personally, I'd recommend adding the Packman repository before doing anything, and then simply executing `zypper dup` from the command line as the root user. This will pull in all the codec libraries from this third-party repository and you'll have pretty decent multimedia support without trying to search and select the codecs manually. Of course, anything that you find missing afterwards can be searched for later.

By the way, if you use one of those Broadcom wireless chipsets, then you can install the `broadcom-wl-kmp-default` (as well as `broadcom-wl`) package from the Packman repository. These are the Linux drivers that were finally released by Broadcom a while back, so there's no need to depend on `ndiswrapper` or `fwcutter` for wireless support any more (if you so choose).

The final point to mention here is that the KDE 4.1.3 version offered is not the vanilla upstream version. It has a handful of features backported from the KDE 4.2 beta branch also, which is indeed nice for us KDE users. However, I'll say

why settle with 'some' backported features only, when kde.org hosts a KDE 4.2 RC Live CD (which I was talking about earlier; and by the time you read this, the final version would be out) on an openSUSE 11.1 base?

4.2: Come to my window

...crawl inside, wait by the light of the moon. Come to my window, I'll be home soon.

Say, what now? Nothing really... I was just trying to impersonate Melissa Etheridge to make KDE 4.2 come to me faster. Seems like it doesn't work this way.

Anyway, on January 14, I headed to home.kde.org/~binner/kde-four-live to grab the Live CD. (Note that I could have configured the 'KDE4-UNSTABLE-desktop' repository on a base openSUSE 11.1 install to pull in the KDE 4.2 RC packages, but what fun is that?) The added advantage of using this CD is that you can also take a look at the progress of KOffice 2.0 and K3b 2.0 with the included beta releases.


To install the Live CD you can follow the same steps as you would when you install the DVD. Now, after installation, I obviously needed OpenOffice.org, in addition to the English UK translations and support, as well as the GNOME desktop. Of course, I could have pulled everything from the online repository. But why waste time when I already had the DVD handy?

Adding the DVD as an install source is simple: *YaST* → *Software Repositories*; in the 'Configured Software Repositories' window click *Add*; select the *DVD* radio button, and follow the steps. That's it! Now, open the YaST package manager and install the rest of the software you need. Simple, isn't it? Of course, by the time you read this article, the final KDE 4.2 will be out, so you should check that out.

Now, I know you expect me to talk about what's new in KDE 4.2, but I won't do that. All I can say is it's a brilliant update with really nice features and a sexier interface. What can I say, other than that I'm loving it!

Bottomline

So, where do we really stand? I'll say, whether you want to try KDE 4.2 or not, it doesn't really matter. What matters is that if you are an openSUSE user, this is a 'must' upgrade. And if you aren't an openSUSE user, well, it's certainly worth a try—at least for the sake of being able to quickly get started with KDE 4.2 without waiting for your distro to release its next version after a few months with the updated desktop.



By the way, the news is that the next version of openSUSE will give v4.2 a skip, and straight away bring forth v4.3, which probably will make it the first distro to offer the new desktop version out-of-the-box this summer. Nice to see at least one major distro trying to follow the KDE release schedule instead of GNOME's, for a change. **END** 

By: Atanu Datta

He likes to head bang and play air guitar in his spare time. Oh, and he's also a part of the LFY Bureau.

getting there together

20 - 21 february 2009. new delhi



freed.in
freedom in technology and software

The **annual event** of the Indian Linux Users Group, Delhi Chapter (**ILUG-D**)
in partnership with Jawaharlal Nehru University (**JNU**)



FOSSTIVAL

Hackers for Hackers

The GNU/Linux Software Development Workshop of GLUG-Meerut became an icebreaker.

*F*OSSTIVAL 2009 (9th-11th January 2009) at IIMT Engineering College, Meerut was like a breath of fresh air within the FOSS world of UP. Jointly organised by 'GNU/Linux User Group-Meerut' and IIMT Engineering College, the workshop focused on grooming budding hackers from all over India. It was one among a few 'FOSS-only' development workshops in north India.

As the GLUG-Meerut coordinator Mohit Singh said, "FOSSTIVAL 2009 is the first step of its kind in the journey of the software freedom movement in UP and the associated regions. As it follows the guideline, 'of the hacker, for the hacker and by the hacker', we firmly believe that it will create a big bang for FOSS in the region and beyond."

Being the first national FOSS event of the

year, FOSSTIVAL 2009 attracted the entire Indian FOSS community by involving hackers of national recognition as well as covering rarely touched topics of wide importance in academics, industry and research. The motive was to act as a catalyst and an appetiser for young minds as well as to attract fresh talent from the large number of engineering colleges in Meerut and beyond.

The workshop had two parallel hands-on tracks from the LAMP stack (called FOSSTIVAL LAMPsticks): 'Open Source Web Development' and 'GNU/Linux System Programming'. The first track focused on Apache, MySQL, PHP and Open Source Projects like 'Sahana', while the second track covered GNU/Linux BASHing, GNU ToolChain, POSIX API, GTK+ Toolkit and Linux Device Driver Development.

A total of 76 engineering students and professionals from all over India participated in the event. The 'breeding hackers' team had Vivek Khurana (Delhi, the guest of honour), Atul Jha (Delhi, *mukt.in*), Gajendra Khanna (student coordinator, IIT-Delhi LUG), Mayank Kumar (student coordinator, IIT-Delhi LUG), Ajay Kumar (Sahana, GSOC '08), Pranava Swaroop (NIT-Jaipur), P. Govind Raj and Jai Bhaskar (both from the BOSS team of CDAC-Noida).

Media partners Radio IIMT-90.4 FM, the first community radio of North India, showed true community spirit by recording the historic promo for FOSSTIVAL 2009. The promo clearly defined a 'hacker' as: "A hacker is not a criminal, but instead, an artistic programmer! Come and become a true hacker at the IIMT FOSSTIVAL workshop." Hats off to the director of 90.4 FM, who took a bold step towards spreading the truth about the 'hacker culture' among the public. The promo has been released under the Creative Commons Share Alike License on www.glug-meerut.org. The director of 90.4 FM, Braham Prakash Yadav, is working on creating an English version

2009

of the promo to be released under the Creative Commons. Richard Stallman wrote the following lines for FOSSTIVAL Promo at Radio IIMT 90.4 FM: "Hacking doesn't mean breaking security, Hacking means playful cleverness. Come join us in hacking!"

Day 1: Booting it up

The first day (January 9) started with a bang! It was an unusual inauguration of a true hacker's event—bursting more than a dozen balloons with names of different non-free software on them! The first balloon was burst by guest of honour Vivek Khurana. This '*Halla Bol*' inauguration thrilled the budding GNU/Linux hackers!

Vivek Khurana took charge of the Web development track by recapitulating the science and technology behind the magic of the Web. He covered the HTTP client and server system, and how they run in a global scenario, the mark-up languages and their role in the Web, W3C, IETF and showed the role of hacking scripting languages for the Web ecosystem. The emphasis was on understanding the protocol standards and development standards for the Web.

Atul Jha initiated the GNU/Linux System Programming track with an overview of the system, shell tips and tricks, as well as a little dose of BASHing. The ground was set for a closer look at the software development ecosystem of the GNU/Linux system using a variety of software tools.

Day 2: Getting the kernel going

The Web track was started by Pranava Swaroop, who demonstrated how to 'configure, learn and run' the grand Apache server. Starting with the internals of Apache and the design of the Apache Web server, Pranava shared some tricks on installing, configuring, hosting and even securing the Apache server. The Apache kernel, *http_core* and libraries excited the young hackers. Other topics of interest with the audience were virtual hosts and *mod_ssl* usage.

Vivek Khurana followed, with an introduction to dynamic languages and the PHP execution model. He then moved on to a hands-on session on PHP. Participants developed PHP applications using the MVC paradigm. Vivek also shared the nuts and bolts of PHP like the data types, sessions, cookies, and OOP. The session ended with an introduction to MySQL and Web services.

In the systems programming track, Gajendra Khanna (IIT Delhi LUG) used the ground set by Atul to interact with GNU/Linux software tools. He started with the famous 'vi versus Emacs' and could not resist interacting with his favourite 'Emacs Psychiatrist' while comparing the two. The Eclipse IDE workhorse was a bonus for the audience. Gajendra demonstrated the intricacies of the GNU Toolchain members using simple programs. He covered GCC, Make, GDB and CVS, which seldom feature in workshops as harmonious elements of a standard toolchain.

Mayank Kumar (IIT Delhi LUG) followed Gajendra by initiating practical IEEE POSIX sessions. He moved from the Linux system call interface and went on to the file system interface, processes and POSIX threads, as well as covering a bit of signals and their usage. The hands-on creation and usage of processes and *Pthreads* gave enough food for thought to budding GNU/Linux systems programmers. The issue of portable applications was discussed in the light of POSIX. Several participants became eager to contribute to the systems programming-oriented open source projects.

Evening 2: Google Summer of Code in the winter of Meerut

In a special evening session, Ajay Kumar (Google Summer of Code winner of 2008) introduced project Sahana—the disaster management system, as a group member of the Sahana core development team. He discussed the brief history of Sahana—the humanitarian FOSS project, the problems faced by it, the technologies used, the deployment scenario, real-life deployments and case studies.

Ajay discussed the Sahana Request Management System, Messaging System and Real-Time Situation Maps with an emphasis on deployment aspects. A special attraction was 'Sahana on Mobiles'.

He inspired young hackers to take up the challenge and opportunities presented by the 'Google Summer of Code'. He discussed the goals of the program, how it benefits the student community at large, what students could do with it and how it benefits them in terms of student stipends.



Nothing like inaugurating a FOSS event by squashing a Windows XP balloon



Busy with the keyboards at one of the workshops



Yes, work smart and y'all shall be freed



Time to take freelunch, the 'shaadi wala khaana'

Day 3: Crunching the CPU

In the Web track, Vivek Khurana shared the nuts and bolts of Web database application development by introducing MySQL to the crowd. He gave a rapid-round recap of SQL and normal forms. Then he introduced a bit of PHP in the MySQL arena. He discussed the benefits of MySQL as a Web database and also helped his audience get a feel of it, hands on.

In the systems programming track, this was the day of the BOSS (Bharat Operating System Solutions) team from CDAC-NOIDA. Jai Bhaskar followed Mayank in an iterative hands-on GTK+ toolkit session, which mesmerised the participants. GUI programming in plain 'C' was an instant hit. Jai combined simple techniques effectively in small GTK+ programs. When the session was over, it was a new beginning for many. He discussed CDAC projects using GTK+ toolkit and inspired the hackers to play with it for a better future for the GUI on UNIX and UNIX-like systems.

P. Govind Raj is a Linux kernel hacker in the BOSS team. He specially prepared a session on 'The Linux Kernel from 30,000 feet', a practical kernel hacking introduction for newbies! His approach worked very well, and budding hackers had a warm feeling towards the kernel and its friends in the GNU system while the temperature outside emulated 'Dilli ki Sardi'.

Govind followed Mayank from the point of the 'syscall' interface in the Linux kernel. He demonstrated how to write a simple character driver and dynamically attach modules to the kernel. He also covered 'strace' and interfacing with 'proc' VFS. The next Indian contributor to the Linux kernel may be someone who attended this session—who knows!

The BOSS team also demonstrated the enhanced multi-lingual capabilities of the latest release, 3.0, of the BOSS distribution. BOSS now supports 18 Indian languages.

The FOSS student projects were discussed largely during the breaks and after the sessions. Sandeep Singh (HBTI Kanpur) was curious to write an open source DoS Defence Mechanism inside the kernel itself. Sumit Srivastava and Navneet Rastogi (IIMT Meerut) discussed enhancements of their LZW compressor and decompressor system developed using GTK+. Arun (MIET Meerut) was curious to leverage the potential of student projects in Web technologies.

Evening 3: Starting the journey

In the closing ceremony, the chairman of the IIMT group of colleges, Yogesh Mohanji Gupta, expressed gratitude towards the historic meeting of national hackers at IIMT Engineering College, Meerut. He also thanked *LFY* for becoming the media partner for FOSSTIVAL 2009, creating another milestone between GLUG-Meerut and *LFY*. All the participants were glad to receive the latest issue of *LFY*.

FOSSTIVAL 2009 was made possible by the collective efforts of Prof. O. N. Mehrotra (director, IIMT Engineering College), Mohit Singh (coordinator, GLUG-Meerut), Arun Singh, Nikhil Gupta, Anshu Sharma, Anjali Shrivastav and the GLUG-Meerut team of the Student Software Freedom Commandos. The FOSSTIVAL 2009 participants left IIMT with a promise to meet again next year in a larger event with a broader spectrum of technology and events. **END** 

By: The GLUG-Meerut Commandos

The commandos helped organise everything from top to bottom and they were also the most enthusiastic participants of FOSSTIVAL sessions. Under the command of Mohit Singh, the GSOC (GLUG Student Onduty Commandos) team of young Software Freedom enthusiasts comprises of Sandeep Singh (HBTI, Kanpur), Sumit Srivastava, Navneet Rastogi, Nikhil Saxena, Fahad Tasleem, Shikha Singh, Farha Naaz, Harshi Gupta (All from IIMT, Meerut), Arun (MIET, Meerut) and many more soldiers of software freedom from all over India. GLUG-Meerut welcomes all software freedom supporters at groups.google.com/group/glug-meerut and is present on web as www.glug-meerut.org



Close

Encounters of the Free Kind!

Three Days with RMS

*T*his was the closest I could get to the man who has changed the world around us by setting us free. Here is my journal of the three days that I spent with Richard M Stallman.

Waiting at the IGI domestic airport on a cold night, I was filled with excitement. My eyes were fixed on the arrivals while awaiting the man we all have admired ever

since we got attracted to the Free Software movement. Richard M Stallman, or RMS as we call him, is a saviour in these times when software 'pirates' are trying ever so hard to take control of our lives—through non-free and proprietary software. You can understand the excitement of meeting a person who has changed the whole world and created a new concept of brotherhood—a community.

Don't allow Big Brother to watch you!

I asked RMS why he doesn't carry a cell phone. He said there were several reasons for that. First, cell phones have a lot of proprietary stuff inside them. Second, he doesn't want 'big brother' keeping an eye on where he is travelling. He also said that since most of the stuff on cell phones is proprietary, we don't know if there are back doors and the companies are gathering information about us.



The reason for being a bit anxious was that RMS doesn't carry a cell phone. Now, I was thirsty and needed water, but I couldn't leave the spot as he might have just passed by and taken a pre-paid taxi. There was no means to contact him and let him know that I was around. I would later ask him why he didn't carry a cell phone. Right then, my eyes searched for him and his friend Dora.

Then I caught a glimpse of the man and I felt an adrenalin rush. He was coming through a glass door, wearing his typical red T-shirt and Dora was by his side. He came out and started walking towards the taxi stand. I yelled "RMS!" He did not hear me. I yelled again and he heard. He came to me. I told him who I was and that I was there to pick him up. I made a call to the cab driver to come and pick us up.

We walked towards the main road. The driver was there and told us that it would be better if we walked to the taxi. He took the trolley from RMS, who hesitated and I said it was okay. At this point he said, "Look, that is the problem with the caste system in India—we are not even allowed to carry our own stuff." He is a critic of the caste system in India and I learned about his concerns other than just free software, which reminded me of *The White Tiger*, Adiga's award-winning book that I had just finished reading. While I tried to talk about several things, he stopped me, saying, "I can't understand what you are saying, please speak slowly." I controlled my excitement, explaining that I was from *LINUX For You* magazine. To which he responded that the name of the magazine is not okay! Well, that is something we can't do anything about now, unfortunately.

We sat in the cab and started towards the guest house at JNU (Jawaharlal Nehru University). On the way he asked me to stop by somewhere and get some water. We found some shops open. While I was about to leave to get the water, he told me, "Make sure its not any of the Coca Cola products because we have boycotted them." Oh! Now, it was already midnight, there was only this one shop open and RMS had put forward this condition. I didn't know anything about the reason behind the boycott, but I decided to keep my questions for later and headed for the shop.

There I learned that they only had Coca Cola's brand of mineral water. I returned empty handed. And that was the second time I came face to face with this man's idealism—he'd rather choose to stay thirsty than drink a

drop of Coke's water. I then asked why he had boycotted the company. He replied that it was because the company used paramilitaries to murder union organizers at Coca Cola bottling plants in Colombia, South America. Now, who ever said that RMS was only concerned about one issue (Linus Torvalds wrote that in one of his recent blogs) proved to be absolutely wrong.

We reached the guest house, but did not find any other shops where we could get water from. At the guest house I offered to try elsewhere, but by then Dora pulled out some small water bottles from her bag and said that they could survive till morning. I bid them a good night and told them I'd pick them up the next morning for a Sun Microsystems' event at IIT Delhi that RMS was going to be a part of. Jaijit Bhattacharya, a good friend and country director, government strategy, Sun Microsystems, was organising the event and had planned RMS's visit to New Delhi.

The next morning we went to IIT where RMS was to inaugurate the Center for Excellence in E-Governance. There he needed to do some mail transfers, and then I realised why sometimes it could take around 24 hours to get his replies. He pulled all the mails from the server onto his small and cute Lemote laptop and then pushed the mails that he had already answered. Then he started answering the mails he had pulled down.

Most of the time he kept working on his laptop, but was always willing to talk and respond to people no matter what they asked. While inside the hall, he noticed the word 'Open Source' associated with the event and got a bit upset. He wanted to talk to Jaijit immediately. I connected them and then he took a promise from Jaijit that they would support the cause of free software. Even during the lamp-lighting ceremony, he said that he was doing this based on a promise from Jaijit that they would support the cause. Ah, I forgot to mention that it was around 12 pm, and RMS and Dora had not had any breakfast. Later, I would learn that he skips breakfast or even lunch if there is work. Well, I've never seen this kind of dedication.

After the inauguration, we had to attend another event—a panel discussion on Policies for Sustainable Absorption of ICT in Society. I got a bit worried to see Windows XP running on the laptop connected to the projector. I told the organisers that it could create some problems and if they wanted they could use my laptop, which runs Ubuntu. They showed their unwillingness to cooperate by saying that it was not possible to switch in the middle of a session. I asked them to hide the Windows XP logo hanging on the screen. I did not want RMS to be photographed speaking against the backdrop of Windows XP.

Someone conveyed that message to Jaijit who was anchoring the show, and they put an obstacle between the lens and the screen. All this time RMS was sitting on a sofa on the stage and reading a book. I thought he was not

interested in the discussion, but I was surprised to see his reactions every time someone used the term 'open source' or 'intellectual property'. Then he took the microphone and clarified certain doubts about the ambiguity of the term 'intellectual property'. He also said, "This is confusion; every time someone used that term, it meant that either the person himself is confused or is trying to confuse others." I was filled with awe at how this man could keep an ear open to what was happening around, while reading a book. Then came his speech, covering in the best possible way the definition of free software, its cause, and the issues surrounding the world of computing.

Later, when Jaijit came to speak he deliberately removed the obstacles hiding the Windows machines. RMS was obviously shocked to notice Windows running there. Jaijit tried to open an ODT file on Windows MS Word and it did not open. The message was clear—Windows can't do a lot of basic things. Jaijit explained that alone was the sole purpose of the Windows machines there. The dramatic presentation left RMS impressed.

Then we went for a meeting, which I can't write about here. But I trust the repercussions of that meeting will be soon known in India.

After finishing with their roles at the event, RMS and Dora wanted to buy some books and clothes, as Dora had not expected it to be so cold in Delhi. We first went to pick up some cookery books that RMS wanted, and then to Dilli Haat to buy some clothes. Now, while Dora was doing all the shopping RMS was either replying to e-mails on his laptop or reading his book. At one point I asked him what he found more difficult—dealing with the opposite sex or with proprietary software companies. He said, "Well, I don't deal with proprietary companies, but I do deal with women... I have no interest in dealing with any proprietary companies, but I do like hanging out with her. But yes, dealing with a woman is a bit difficult."

I recalled a discussion with one of my colleagues who once said that a country like India needed dictatorship after independence to first teach people how to be good citizens before handing them the power of democracy. I asked RMS what he thought about that. He felt that democracy was always the best option, because the one who has power may abuse it. But in a democracy you have the choice to replace the person. His reasoning was based on the question: "What is the guarantee that the dictator you bring in will do the right things only?" Then you will have no power to replace him. And things will take their own sweet time to improve.

After that we bought some more non-Coca Cola water and went for a dinner organised by Jaijit. Over dinner he told me that he also writes science fiction and booted his laptop to show me two of his science fiction stories. I was just amazed at this man. He has given us free software, he has created the GPL, and he is a sci-fiction writer as well! The stories were awesome and I became a fan immediately. Even during the dinner, which he enjoyed a

A preacher who practices


While on the move in Delhi, we were approached by some beggars with a wounded child. While Dora gave out money, RMS refused. He said that the kid may not be really wounded, and said that they just find it easier to beg. There was a minor argument between them. He was clear that if somebody was asking for money for food, he would rather go and buy the food for them. Interestingly, at Janpath, a beggar came to us with a Euro coin. RMS took the coin and started calculating how much that meant in Indian currency, and then pulled the money out and handed it to the beggars. Well, he sure is a guy who not only preaches but also practices what he preaches!



lot, he would pick up his laptop and reply to some e-mail or the other. It seemed he wasn't in a mood for (free as in...) free beer like us, and settled with some other drink. After that we returned to the guest house. He again needed the Internet connection, which wasn't available in his room and I looked for some browsing centres while they waited in the car, but everything was closed, as it was already midnight.

The next day, RMS was going to deliver a lecture at the JNU. In the morning, we first met a few government secretaries for media, IT and science. There he was convincing them about the benefits of using free software and why it mattered to governments. We came back to JNU, and then Dora needed to do some shopping so I took her to the Janpath market (near Connaught Place in central Delhi). The next day they planned to go to Agra. And I went back to my office.

RMS was to leave Delhi on Sunday, so we went out for lunch. It was nice food—though, he didn't like it spicy. I shared with him the science fiction I was writing on free software; he listened to the whole story and said it had potential, before making a lot of suggestions as well as helping me tighten the plot a bit. We then went to the airport and he left for Hyderabad.

I don't know if this was a dream or reality. I got to spend three days with the man who has changed the shape of this world and has ensured that we will never become slaves to the multinationals. He has set us free, I thought, as I bid goodbye to the free man! **END** 

[Note: This is my interpretation/perception of RMS. If there are any minor differences in his actual idealism and my interpretation of it, the mistake would be mine.]

By: Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan, Swapnil also writes fiction and tries to find cracks in a proprietary company's 'paper armours'. He is a big movie buff, and prefers listening to music at such loud volumes that he's gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of *EFYTimes.com*.

A close-up portrait of Linus Torvalds, the creator of Linux. He has long, wavy, greyish-brown hair and a full, greyish-brown beard. He is smiling slightly, showing his teeth. He is wearing a light-colored, possibly white, shirt. The background is out of focus, showing green foliage and a small, light-colored rectangular object with a picture on it.

Face to Face with **RMS**

It was one of the rarest opportunities to spend some time with the man I've always looked up to since I migrated to GNU/Linux systems. The father of the free software movement, Richard M. Stallman was in Delhi for three days, and during the time I spent with RMS, we discussed various topics that helped me get to the bottom of why 'Free Software' matters so much. This frank discussion with RMS also covers topics like the entire confusion created around patents and copyright, what a sensible government should do, why free software matters for an enterprise as well as a developer and what will happen to the free software movement after RMS... though we wish him to live another 100 years. So read on...

Chilly Pepper: This is the first interview in which RMS has responded to a recent blog post by Linus Torvalds, where the creator of the Linux kernel has criticised the father of the free software movement.

Q I would like to start off with what free software is all about, according to you?

Free software means software that respects the user's freedom and the user community. Proprietary software traps the user's freedom and divides the user, leaving him helpless. They are divided because they are forbidden to share the source code. Helpless because they don't get the source code; hence they can't change the program and they can't even verify what it is doing to them. Free software respects that and there are four essential freedoms that a user must have.

Zero is the freedom to run the program as you wish. Freedom One is the freedom to study the source code of the program and then change it to make the program do what you wish. Freedom Two is to help your neighbour, which means to make exact copies of the program and distribute it to others when you wish. And Freedom Three is the freedom to contribute to your community, which means distribute modified versions when you wish.

Now, this has nothing to do with the details of what the program does and how. It's about what you are allowed to do with the program. So, if one of these freedoms is missing, or partly missing, then the social system of a distribution is unethical, and that makes it proprietary software. So, that software should not exist, because every time someone uses it, it becomes a social problem.

Q To whom does free software matter?

It matters to everybody who uses software. So, if you have a computer these freedoms are important to you and you deserve free software.

Q Is there any particular reason why free software makes more sense in a country like India?

I don't think so. I think freedom is essential for everyone. The only country where free software doesn't matter is the country where nobody uses computers.

Q Why is it important for governments?

All users of software deserve the four freedoms and that includes government agencies. However, in their case it's not just what they deserve, it's their duty, it's their responsibility. Because the government has a responsibility to maintain its sovereignty, i.e., control over what they do. So a government agency that does computing has a duty to maintain control over its computing, which means it must use free software.

Q What are the kinds of relations that exist between the governments and the free software community around the globe?

It depends on which government—they are not all doing the same things. Some governments have explicitly adopted a policy of moving government agencies to free software. These include Venezuela, Ecuador, apart from others who are migrating to free software though they don't have a formal policy.

Then there are some governments like the US government that has no pride, no ethics, and basically just sells out to multinational corporations. And then there are some governments that are very concerned about letting those corporations have power and try to resist it. And while doing so they also defend their citizens from those corporations.

Q And what about education? Now there are two aspects to that: using free software in the IT infrastructure of schools, as well as getting it into the curriculum.

Well, regarding the administration of the school, they deserve to use free software just like everyone else. What is especially important is what software they teach students. Because schools have influence over society and they have a social mission—which is to teach the next generation to be good citizens of a strong, capable, independent, cooperating and free society. In the field of computing, this means that their responsibility is to teach students to be users of free software and to appreciate and applaud the freedoms of free software.

If a school is teaching proprietary programs, that means it is selling dependency to the next generation, directing them under the power of a particular company. This is something schools must never do. In addition, some students are people who will be fascinated with computers and will learn to be great programmers. When they get interested in programming, which is usually between the age group of 10 and 13, they want to learn how the software

works, and then they want to start changing and improving it. That is how they learn to be great programmers. So, if the school uses free software, the students have the opportunity to learn. If the school uses proprietary software, children don't have the opportunity to learn, because proprietary software is an enemy of education, it's opposed to the spirit of education. And thus, it doesn't belong in a school.

Q Do you think the training cost is one of the concerns? If you look at the education system, it needs investments to train teachers on the new systems?

Well yes, in order to get out of a path that is leading you to a bad destination you have to make an effort. And this is a kind of investment, social investment. Even if you ignore all the freedom and the community, the cost savings alone are worth making that investment. That is the smallest and most superficial part of why you should do it.

However, you will notice that the proponents of proprietary software encourage people to think short term, because if you think very short term it might seem that switching to free software is a lot of work and more expensive. It seems more expensive because you have to train some people at first. But that is a one-time expense. Once schools start graduating people who know free software there is no shortage of people. So, this is only a short-term expense and, therefore, those who are trying to defend proprietary software are trying to encourage a very short-term way of thinking.

Q How much sense does free software make to the enterprise segment?

Well, if you are running a business, you ought to want to have control over your computing, so you should insist on free software. Businesses get the benefits of the four freedoms just like everyone else. Not only that, there is a side effect. The support for free software is a free market. In case of proprietary software, only the developer has the source code, so if there is a bug, only the developer can fix it and thereby support is a monopoly. This means that support is usually expensive and bad. But with free software that support is available from the free market, so businesses get better support for their money. And any time they are not satisfied with their support provider they can switch to another.

Q What about developers who while using a lot of free software in the backend, release their products as proprietary software?

Well, that is unjust. Any proprietary program means that there is a user who doesn't have freedom. There is no excuse for this. Users should refuse to get subjugated in this way. They accept it because of the short-term thinking—"I want to do this now, and the only way I can do this now is to let this developer have power over me." So, what is to be done? Well, that depends on whether I am thinking short-term or long-term.

The free software movement has made it even easier to

choose freedom. When I started the movement 25 years ago, the only way I could have freedom was through developing an operating system and that was going to take years. If I were thinking short term, I would have said: "Oh, it's so far ahead. I give up!", and I would have let the developer of the proprietary software have power over me the same way everyone else was doing. But instead I thought long term, and I thought: "Yes, it would take years, but it is worth it!" Today, we have free operating systems, we have thousands of free applications, and as a result, choosing freedom doesn't require you to work for many years any more. Now, at most, it requires an inconvenience every so often.

Q But why should developers be working on free software?

I am not saying that everybody should develop free software. If you don't feel like developing free software then you don't have to, because lots of other people will. I certainly don't want to pressure anyone into spending his time writing free software if he is not interested in doing so. But what you should not do is release proprietary software, because that is denying freedom to somebody else. It is not legitimate and is unacceptable for anyone to do that. So, if you are developing software to release, you must make it free software. But you don't have to develop software to release it.

Now, we should know that most people who are paid to write software are writing custom software. There is a particular client that wants a particular program, and he is paying them to write it. This client should, of course, insist on receiving the product they paid for as free software and that should be capable of running on a completely free platform. Because, otherwise, they are paying to lose their freedom. They still have to pay some people to write solutions, because this is not something people would want to volunteer for otherwise. So, companies have to pay and this means that most people who get paid to program free software won't change things much. It may increase their productivity somewhat, but it won't change things much.

Q There's been a lot of confusion between free software and open source. Would you like to say something about it?

Once Linux was put together with the GNU operating system in 1992 to make a complete free operating system, people started distributing the GNU/Linux operating system and telling friends about it, and they were mainly techies. People accustomed to judging software in technical terms, recommended the system to other people looking at the practical advantage of GNU/Linux systems and they did not mention these ethical issues.

Meanwhile, we in the free software movement were talking about these ethical issues. Some people listened to us and others listened to them. So during the 90s, a split, an argument, arose within the community between the people who valued the freedom and community above all and the people who only valued practical convenience.

In 1998 this group, which valued only superficial practical values, chose the term open source as a way

“If someone starts making statements using ‘intellectual property’, that is a sign of confusion, because that term attempts to generalise several different laws (copyright is one of them), when all these laws have nothing in common. So, whatever that person says is pure confusion; it is nonsense. And this nonsense is being passed off as meaningful.”



to avoid mentioning or alluding to our ideas of freedom. Because that term had never been used, they could choose whatever idea to associate with it and leave out other ideas they did not want to mention. That's what they did. Ever since then, free software and open source are two fundamentally different philosophies based on different ideas of what is important.

So, the difference between the two philosophies is not just the detail, but goes down to the root. We have some criterion for free software, which is actually the criterion for the licensing and distribution of the program. There is also a definition of open source that has a criterion of how a program is licensed and distributed. So, in practical terms, they come out fairly similar. As far as I know, all free software is open source and nearly all open source programs are free software. But there are some open source programs that are non-free, and the reason is that they interpret their criterion for licensing in a way that is little bit more lax. So, there are some licences that they have accepted, but we have not.

Q The ordinary public uses the word open source...

Well, that depends... Don't simply assume that, and don't declare the total defeat of the free software movement, because that is not true. You just made a statement that is not true, because the ordinary public is a lot of people. Some have heard of one and some have heard of the other. It's true that more people have heard of open source, but that doesn't mean that everybody has heard of only open source.

The reason is that the supporters of open source are more in number and thus the companies who are involved with software mostly say 'open source'. And the reason is that most of those companies also are involved with proprietary software. They don't want to educate the public to reject the proprietary software on moral grounds, thus avoiding mentioning the free software movement and never mention our ethical ideas. They can win certain amount of favourable public opinion by connecting themselves with open source, and yet avoid teaching users to reject proprietary software.

They find a way to use the term open source and

their company's name in the same sentence even if what they are saying has nothing to do with it. They don't generally do that with free software because the idea of free software says that if it doesn't respect freedom it is an injustice, and so you should reject it. A company that makes non-free software is not going to want to link its name with free software. And that's why, theoretically speaking, they could play these double-talk terms with any other terms, but they don't generally play these tricks with free software. They do a different thing, however—they misunderstand the term 'free' and interpret it as gratis copies of proprietary software telling that it is free.

The English language has a flaw that other languages don't have—that is, there is no word that means free as in freedom with only that meaning. Whereas in most other languages there is a clearer word which doesn't mean zero price. In France, I always say 'logiciel libre', because 'libre' means free as in freedom. So, most languages have a word which mean free as in freedom and it doesn't talk about price. English seems to be the only language that fails to have such a word.

Q Can you clarify a bit between copyright and patents?

First thing, copyright and patents have nothing to do with each other. So, you may as well ask me to clarify between copyright and trees... or side walks.

Q But at times, people do try to club them together...

The first thing to realise is that copyright and patents have nothing in common; they are totally unrelated laws. What one of them does, the other one doesn't. So, in particular, if someone starts making statements using 'intellectual property', that is a sign of confusion, because that term attempts to generalise several different laws (copyright is one of them), when all these laws have nothing in common.

So, whatever that person says is pure confusion; it is nonsense. And this nonsense is being passed off as meaningful. But, most people don't know it is nonsense, so they repeat the nonsense. There are also people who

know how wrong it is, but they use it and then they justify it. There are a lot of others who are simply trying to mislead other people. It is in their interest not to have other people understand these issues.

Q Is there any connection between the two at all?

What has copyright to do with licensing and what does it have to do with free software? Well, under today's copyright law, any work that is written or composed somehow is automatically copyrighted. So, there is a copyright on that work or particular authorship. By default, the copyright law says that people are not allowed to copy or distribute or modify the program. In some countries people are not allowed to run it without permission. So, how do you make free software?

The only way is through a formal declaration by the copyright holder saying that you, the user, have the four freedoms. That declaration, we call the free software licence. It is not the only context in which the term 'licence' is used. Licences are signed contracts, but this one is not. This is simply a unilateral grant of permission from the copyright holders of the program. Free software licences have some conditions. It could be a tiny requirement like you can't remove this licence, but some other licences have more requirements. For example, I wrote a licence, which is one of the many free software licences, called the GNU General Public Licence or GNU GPL. And it is used in over two-thirds of all the free software packages. However, GPL is not all—there are other free software licences too.

What's special about this licence is there are certain requirements we call copy-left and that is: "When you redistribute copies of the program, either exact or modified, you must keep the licence the same and you must distribute the whole of the modified version under the same licence and make the source code available." These licences make an extra effort to defend the freedom of every user.

Q What led to the move from GPL version 2 to 3?

That is hard to say, because there were different

reasons for every change. If you look at GPLv3.fsf.org, you can find the rationale document that states the purpose of every change that we made. However, in general, the overall purpose of version 3 of the GPL is the same as version 2—to defend the freedom of every user.

We made some changes for convenience. We made some changes in order to permit things that ought to be permitted. And we made some changes to defend users' freedom against new methods of taking it away. For instance, there are some changes we made to better defend freedom against software patents; we made some changes to block tivoization. We also made some changes so that you can lawfully redistribute the software using BitTorrent. Because it turns out that BitTorrent violates GPL v2. But, obviously, it is a good thing for people to use BitTorrent as long as they use it for the right purpose. So, in GPLv3 we added a clause to make it legal.

Q What do you think about the MS-Novell deal?

The deal between Microsoft and Novell was an attempt to take away our freedom. It is an attempt to convert free software into proprietary software. The way they did it is by using Microsoft's patent claims. A patent is a government-issued monopoly on an idea. Some countries allow patents in the software fields. When that happens, it means that a software technique or feature can be patented. If a country allows software patents then all the developers are in danger, because when you develop a program, you implement many ideas. If some of these ideas are patented that means the patent holder can sue you. If you distribute the copies to some users, then the patent holder can sue those users too in some countries.

The mega corporations own half of the patents; that's true in every field and they cross-licence each other in a way that gives them a form of dominion. This is why the software mega corporations want software patents. Three years ago, the mega corporations in the US asked the government of India to implement software patents and the government went along with it, but the opposition parties blocked it.

"You can see that he [Linus Torvalds] is a person who doesn't believe in freedom. You can tell that from his writings. Why does he reject GPLv3, because GPLv3 protects users from tivoization, which is a fairly new method of denying freedom. It did not exist when I wrote GPLv2 or I would have done something about it then."



So, they were not able to change the law. Now the Indian patent office is trying to authorise the software patents again by twisting the existing law. So, you really have to fight constantly to stop these companies from grabbing control.

Now, in the MS-Novell pact, Novell accepted a licence patent for Novell customers from Microsoft and Novell agreed to pay Microsoft a fee for usage. So, if you are a Novell customer you are paying Microsoft. Because of this, there is a boycott of Novell by free software activists who condemn this betrayal.

Q What do you say about censorship?

I should explain that this is not part of the free software issue. It's a much older issue of human rights. Censorship is an evil. It attacks your freedom and is very dangerous.

Q So, you mean there should not be any kind of censorship?

No, not at all. Now, I have seen things that disgusted me. For instance there was a movie, which was very famous and very popular, called *Pulp Fiction*. The violence in the movie disgusted me. If there is somebody bleeding on the street and I can help him, I guess I have to force myself to look. But if it is just fiction in a movie, why should I suffer?

However, I am against censorship. I do not want people to censor movies like *Pulp Fiction*. I am against it. I defend their freedom to make movies like these, and I am not going to go and watch them.

Q There has been a lot of concern regarding software as a service. What do you think?

Software as a service means that in order to do your computing, you are going to send your data to someone else's computer. There, he is going to run a program and your computing will be done in his computer and with his copy of a program. If the program is someone else's copy then the person who owns the copy takes away the control from you. Here you don't have control on your computing, so you must not do this. The only way you can have control of your computing is if you do it in your copy of a free program. But the structure of the problem is different.

Interestingly, in case of software as a service, it doesn't matter whether that guy's copy, which you are using, is free software or not, because if it is free software—that means he has control of the copy and not you. We can't say that you, we and everyone has to have control over his copy. Because he is entitled to control his copy. That means, when we do computing we should not use his copy; we should use our copy. That is the only way the user can have control. So, software as a service is simply bad and cannot be redeemed. Nothing can make it acceptable.

However, I want to point out that most servers are not software as a service. Most servers are just doing something and, in most cases, it's not you computing your data on that server. For example, if you visit *stallman.org*, you can read what I have published. That is not

software as a service. Or if you go to the free software directory at *directory.fsf.org*, you can search our list of free programs. Well, that runs a program, but that is not software as a service. What you are doing is looking at our data. So there is no problem. Or suppose you go to a Web-store and buy something; what you are doing there is e-commerce and that is a mutual transaction between you and someone else. It is not you computing your data, so that is not software as a service. Suppose you help edit a Wikipedia page—what you are doing there is helping to do Wikipedia's computing. And that is not your computing on your data. It is not software as a service and it is fine. So, we should be aware when I condemn software as a service—out of all the websites in the world software as a service is a tiny fraction of them. It is a narrow way of using a server. Most Web servers are something else.

Q But when I am running a website, my data is on someone else's server and I don't even know where my data is.

That is a somewhat different issue. You can have your computer that is hosted by somebody and you can still have your copies of software. So, you still do control the computing that is done. True, you are trusting your data to whoever hosts the machine. He can go to the machine and rip off the data. It is trusting somebody, but at least he doesn't control the computing that you do.

Q But today people care about mobility; they don't want to carry their data. They want to access it from where ever they want.

That is silly! That means you are trusting your private data to someone you don't know you can trust. So, that is foolish, that is really dumb!

Q There shouldn't be any such thing?

I don't think you should do it. In many cases it goes hand-in-hand with software as a service. Consider the worst example, Google Docs. With Google Docs here is what happens: you keep your private data, or you wish it were private, on Google servers. So, you are trusting a multinational corporation with your privacy, which, of course, you can never trust—any of them for that matter. In order to use it, the site transmits a large proprietary program to your computer, which you may not realise as it happens invisibly – it's a Java script program and it gets loaded into your browser.

Aside from this, which is obviously wrong, there is also a large program that is an unreleased program running in Google servers, and Google, of course, controls it and you don't. So, this is absurd—you should have control of this computing. You should run your word processor or your spreadsheet on your computer. Don't run it on Google's computer or anyone else's. And I am not saying that Google is particularly worse, although it's true that they are. They should release that Java script

program—that will solve one of the problems.

But the other problem is the program in the server; there is no solution for that. Even if Google releases that as free software, it would not solve this problem. The fact is, if you are using someone else's server, no matter how ethical that someone is, no matter how much you can trust him, you still don't control your computing.

Q What do you say about DRM?

You mean digital restriction management—digital handcuffs? This refers to more malicious features put into some programs to restrict the user, most often to restrict your access to your copies of published works. This is injustice! DRM must not exist. And almost always DRM is imposed by a conspiracy of companies. There are a few exceptions, but they are minor.

Consider, for instance, the inscription on a DVD. That scheme was set up by a group of companies. They established the DVD conspiracy that anybody who wants to make a DVD player has to join them and promise to design DVD players so that they restrict the users, just like every other DVD player manufacturer. Now, I believe that these conspiracies ought to be felonies. So the executives that organised these conspiracies and signed up to participate in it should be prosecuted and imprisoned. But, unfortunately, our governments are sell-outs and they have taken the side of those conspiracies, against the public. These executives are well aware that they are in no danger of being prosecuted.

In fact, many governments have passed laws forbidding us to get around the restrictions they have imposed on us, starting with the US. For instance, there is a free software that can play a DVD and in the US that free software is censored. This shows how evil they are. They actually practice censorship of software. You may ask how I know about these conspiracies. It's because it's no secret—they boast about these conspiracies. They are so confident that they will not be prosecuted for this attack on our freedom and this conspiracy to restrict commerce that they stand upright and say, "We have made this agreement to restrict the public."

So, DRM is evil. We should establish laws to prosecute conspiracies that impose DRM, and we should completely reject all products that have DRM unless we have the free software that can break the digital handcuffs. Never buy an encrypted DVD, or any other product with DRM, unless you personally have the means to break the DRM. If you have a free software program to watch the encrypted DVD, then it is OK to buy an encrypted DVD. But otherwise you must only get unencrypted DVDs. In India, I guess the legal DVDs are encrypted so you should not buy them, and I think the illegal ones may not be encrypted so those are OK.

Q I am excited to know about the status of development of the Hurd?

The Hurd is the GNU kernel that we started in 1990

because there was no free kernel and we needed a kernel to have a complete GNU system. Well, various things went wrong and it works, but it doesn't work all that well. However, it's making progress. A few years ago they could only compile 40 per cent of Debian packages on Debian GNU/Hurd and now they can compile 60 per cent. So, they are making progress, but have a long way to go.

Q Do you have any idea when it will be ready?

No, I have learned not to make predictions about such things. There is a lot of work to be done and progress is being made by volunteers. Now, we don't consider finishing the Hurd as one of the high-priority projects. Although we would like to have Hurd work, but the fact is there are free kernels. So, it is not an area where the community has nothing. The high-priority projects are to do things that free software doesn't do.

Q Linus Torvalds recently criticised you on his blog at torvalds-family.blogspot.com/2008/11/black-and-white.html. Do you have something to say about that?

Well, you can see that he is a person who doesn't believe in freedom. You can tell that from his writings. Why does he reject GPLv3, because GPLv3 protects users from tivoization, which is a fairly new method of denying freedom. It did not exist when I wrote GPLv2 or I would have done something about it then.

Tivoization means delivering a product with some free software, but the machines are designed so that if the user installs a modified version, it won't run at all. So theoretically, the users are free to study and change the source code, but they can't run their own binaries. This means that if they change the source code, it's useless. So practically, they don't have the freedom to study and change the program. Therefore, I decided in GPLv3 to prevent this. V3 requires the manufacturer to provide to the user the information necessary for the user to install his own changed binaries in the product he owns and make it run, assuming it is possible—if it's denied by law, then it's okay, because no one can change that.

We are trying to stop a practice where a manufacturer can change it, but the user can't. Well, this is what Torvalds objects to. He is in favour of tivoization. He doesn't care if the user of, in this case Linux, is free to change it. So, what can we do? He has a right to his views. I am not a one-issue person and I care about a lot of political issues, as you can see if you look at stallman.org. But the free software movement is a one-issue movement. Lots of people support the free software movement, who have different views on other issues. That's why I put my views on other issues into my personal site—they are not part of the free software movement.

Q This whole free software movement has sort of grown bigger than RMS...

I hope so, as I am not going to be around forever.

RMS visits a computer centre for the poor

Ambedkar Community Computer Centre (AC3) is situated in a slum in New Gurupampalaya in Bannerghatta Road, Bangalore. In the IT city of multinationals, it is only one of the many slums. On December 13, 2008, this slum had a unique visitor, famous the world over as the founder of the Free Software Foundation. Richard M Stallman's visit turned the attention of IT professionals and media persons to this otherwise 'invisible' slum.

AC3 was set up by members of Stree Jagruthi Samithi, volunteers from AID (Association for India's Development) and by some Free Software Volunteers working in the IT industry.

The idea of setting up the centre was to equip the poor children in the nearby area with the basic computer skills using quality software, which is shareable and available free of cost. The centre uses free software tools on GNU/Linux to teach and impart computer knowledge to the children. The effort is now showing results and the children have shown great enthusiasm. One of the children, Mani, who studies in Class 9, has created beautiful pictures using The GIMP. His pictures got a huge response from the exhibition conducted during the Free Software National Conference held in Cochin in November 2008. Amazed by Mani's pictures, RMS agreed to visit AC3 and interact with the children for a few minutes.

During his visit, RMS inaugurated Mani's painting exhibition and released a book written by children on the basis of their experience in this computing centre. Sarasu, who is from the slum area and teaches kids at the computing centre, presented a talk about AC3 and the work it is doing for the poor. RMS made two corrections to the presentation. He emphasised that free software is not the alternative, but it is *the* right one. He also added that instead of saying "a few success stories" the presentation should say "a great success story" as this community computing centre is doing a great job.

There were also some cultural events presented by the slum dwellers, after which, RMS interacted with the audience and addressed their questions. Here is an extract from Stallman's speech:

"I am honoured to be here in the community computer centre that is in the name of Dr B R Ambedkar. I have read



about Ambedkar and I was inspired by the work he has done for the dalit people in India. There are many issues among the poor and marginalised in India that are more important than using free software. But, free software is one of the tools that helps the poor and the dalits to resolve some of the issues they are facing. Poor people can't afford proprietary software and free software helps them get access to computers. While proprietary software companies are like the colonial rulers who exploited the masses, free software is freedom and liberation. The Ambedkar Community Computer Centre is a model for the other parts of India in the way it gives access to computers to the poor without compromising their freedom. The Ambedkar Community Centre is doing a great job by empowering the poor with knowledge of free software."

It was a memorable experience for the people in the slum. They were all inspired by the visit of RMS to spread the message of free software and also take the centre to the next level.

By: Sreejith G.S.

The author is a free software supporter and is part of the Ambedkar Community Computer Centre team. He works as a technical writer at Sun Microsystems.


Q Right, so how self-sustaining is it?

Well, your guess is as good as mine. But I can see there are thousands of free software activists. There are other people who could tell you the same things that I can tell you, but there are people who would listen when I say it and they perhaps won't when other people say the same thing, because my name is better known now. Well, we make use of that for the better cause.

Q What are you doing to ensure that your movement is self sustaining?

Well, I do my best. Basically, what is going to increase our chances? Having more people who are free software activists and who look for ways to be better activists—

who look for ways to spread the idea of the movement and try to learn to do it the best they can. That's how they will be the same sort of person that I am now.

Well, I hope the force of free software activists keeps growing...  **END**

Interview by Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan, Swapnil also writes fiction and tries to find cracks in a proprietary company's 'paper armours'. He is a big movie buff, and prefers listening to music at such loud volumes that he's gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of *EFYTimes.com*.

Recipes for Networking



It's always fun to try out different hacks under the GNU/Linux freedom platform. The pride of becoming a command-line wizard makes everyone stay close to the CLI. Moreover, the CLI vests you with the ultimate power to control your machine.

In this article, we'll explore networking under GNU/Linux. You'll find it interesting to manage the entire network through certain valid keystrokes known as commands. Imagine that you have to access the contents of several other machines from a mount point in your machine. Then imagine shutting down, rebooting, and installing applications on those remote machines, all at one time? Could you configure the WLAN and LAN interfaces from the CLI? This tutorial gives you some insights to the exciting bytes on controlling your network under GNU/Linux.

First we will learn the 'Hello World' of a networked machine.

Let us ping!

ping is a universal command that is available on every operating system to test the reachability of a network. When you shoot your terminal with the *ping* command and an IP address as its argument, the machine will try to send some bits of raw data towards the machine owning that IP address. If some machine exists with that IP address, it will send back certain bits. Thus the machine receives the bits and it confirms that a path is available from the current machine to the other through a network. We can check the existence of certain machines on the network by just ping.

To see which machines are up in the current network, let's write a bit of shell script. Open *vim* as the root:

```
# vim /usr/bin/netup.sh
```

...and key in the following lines in it:

```
#!/bin/bash
for i in 192.168.1.{1..255}; // checks 192.168.1.x class of IPs.
do
ping -c2 $i > /dev/null;
[ $? -eq 0 ] && echo $i is up.

trap "exit" SIGINT // To force exit when Ctrl+C keystroke is
applied.
done
```

Save the file, and make it executable by running the following command:

```
chmod a+x /usr/bin/netup.sh
```

Now, run the script as:

```
[slynux@gnuxbox ~]$ netup.sh
192.168.1.1 is up.
192.168.1.3 is up.
192.168.1.4 is up.
```

Configuring your network

Now, let us look at how to configure your machine on the network. You can configure it using two methods. It can be configured manually by the *ifconfig* command for static IP addressing or it can be done via the DHCP (Dynamic Host Control Protocol).

Static IP addressing is the one that you explicitly instruct the system to use by giving an IP address for a given Ethernet or wireless interface. In case you're using the DHCP, simply issuing the *dhclient* command will fetch the

system an available IP address from the DHCP server in your network. Note that it may not be the same IP address that your machine fetches each time you issue *dhclient*.

Interface cards

Machines are networked either via network cables or using wireless protocols. LAN cards used for networking are known as Ethernet and wireless LAN (WLAN) cards. We interface the network via this outlet. In *nix platforms, Ethernet cards or WLAN cards are denoted as *eth0*, *eth1*, etc, or *wlan0*, *wlan1*, etc, respectively.

ifconfig

We have *ifconfig*, a.k.a the interface config, for setting up a network on the machine. To get information about the availability of interface devices available on the current machine, open a terminal and execute the following as the root:

```
[slynux@gnubox ~]# ifconfig -a
eth0    Link encap:Ethernet  HWaddr 00:1C:23:FB:37:22
        inet6 addr: fe80::21c:23ff:feb3:722/64 Scope:Link
        UP BROADCAST MULTICAST  MTU:1500  Metric:1
        RX packets:9724 errors:0 dropped:0 overruns:0 frame:0
        TX packets:2720 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:2400589 (2.2 MiB)  TX bytes:645396 (630.2 KiB)
        Interrupt:17

lo       Link encap:Local Loopback
        inet addr:127.0.0.1  Mask:255.0.0.0
        inet6 addr: ::1/128 Scope:Host
        UP LOOPBACK RUNNING  MTU:16436  Metric:1
        RX packets:76320 errors:0 dropped:0 overruns:0 frame:0
        TX packets:76320 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:87151068 (83.1 MiB)  TX bytes:87151068 (83.1 MiB)

wlan0    Link encap:Ethernet  HWaddr 00:1C:BF:87:25:D2
        inet addr:192.168.1.143 Bcast:192.168.1.255 Mask:255.255.255.0
        inet6 addr: fe80::21c:bfff:fe87:25d2/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:45302 errors:0 dropped:0 overruns:0 frame:0
        TX packets:37510 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:31091293 (29.6 MiB)  TX bytes:9734025 (9.2 MiB)
```

Here I have three interfaces -- *eth0*, *lo* and *wlan0*, where:

- *eth0* corresponds to the Ethernet card
- *lo* corresponds to a loopback device that points to the localhost network
- *wlan0* corresponds to the wireless LAN card

Static IP addressing

For static IP addressing, issue the following command as the root:

```
ifconfig <device name> <ip address>
```

For example:

```
ifconfig eth0 192.168.0.2
```

ifconfig -a gives you details of all interface devices and configurations. In order to receive details of only one Ethernet device, execute *ifconfig eth0*. The following is an example output:

```
[root@gnubox slynux]# ifconfig eth0
eth0    Link encap:Ethernet  HWaddr 00:1C:23:FB:37:22
        inet addr:192.168.0.2 Bcast:192.168.0.255 Mask:255.255.255.0
        inet6 addr: fe80::21c:23ff:feb3:722/64 Scope:Link
        UP BROADCAST MULTICAST  MTU:1500  Metric:1
        RX packets:9724 errors:0 dropped:0 overruns:0 frame:0
        TX packets:2720 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:2400589 (2.2 MiB)  TX bytes:645396 (630.2 KiB)
        Interrupt:17
```

Now, let us go about setting the subnet mask. This is done easily with *ifconfig*, as follows:

```
ifconfig eth0 192.168.0.2 netmask 255.255.255.0
```

Now that the IP address and subnet mask is configured (or reconfigured), how do you get your Ethernet up (i.e., available) and down? The *ifup* and *ifdown* commands help you with that as follows:

```
ifup eth0
```

```
ifdown eth0
```

Wireless networking

In order to hack a wireless card, we have another utility called *iwconfig*. It works similar to *ifconfig*, but it has lots of additional features that are bonded to wireless cards. If we are using a wireless network with static IP, we can attach our wireless card interface to a network as follows:

```
iwconfig wlan0 essid slynux
```

...OR:

```
iwconfig wlan0 essid slynux key 8c140b2037
```

...where 'slynux' is the ESSID (that is, the name of the wireless network) and '8c140b2037' is the security key. Of course, you need to replace these variables with the values that hold good in your network. You can also scan and check the availability of wireless network(s) in your vicinity using the *iwlist* command as follows:

```
[root@gnubox~]# iwlist wlan0 scan
wlan0 Scan completed :

    Cell 01 - Address: 00:08:5C:52:E9:83
        ESSID:"slynux"
        Mode:Master
        Channel:11
        Frequency:2.462 GHz (Channel 11)
        Quality=92/100 Signal level:-39 dBm Noise level=-78 dBm
        Encryption key:off
        Bit Rates:1 Mb/s; 2 Mb/s; 5.5 Mb/s; 11 Mb/s; 18 Mb/s
            24 Mb/s; 36 Mb/s; 54 Mb/s; 6 Mb/s; 9 Mb/s
            12 Mb/s; 48 Mb/s
        Extra:tsf=00000000fc021187
```

The above command will list out the various wireless networks available with a number of properties.

Then we can set the IP for the interface card using the *ifconfig* command itself:

```
ifconfig wlan0 192.168.0.5
```

If you are using dynamic addressing, you can obtain the IP address as follows:

```
dhclient wlan0
```

The settings that you've configured with the *ifconfig* tool are available until the system reboot. But it's a waste of time if you need to configure it on every system start. And hence we take the aid of network configuration scripts. On an Ubuntu (or any other Debian-based) system, this file is located at */etc/networks/interfaces*, and contains data similar to the following:

```
auto lo
iface lo inet loopback

iface eth0 inet static
address 164.164.32.101
netmask 255.255.255.240
gateway 164.164.32.97
```

It is necessary to learn this scripting in order to play with your network. The syntax for these are as follows:

1. Add the following lines if you want to configure eth0 as the DHCP:

```
auto eth0
iface eth0 inet dhcp
```

2. Add the following files if you want to configure static IP:

```
auto eth0
iface eth0 inet static
```

```
address <ip_address>
netmask <netmask>
gateway <gateway_ip>
```

3. If it is a wireless network, add the following lines along with the above lines:

```
wireless-essid <network_name>
wireless-key <key>
```

Now, to restart the network daemon, execute the following as the root:

```
/etc/init.d/network restart
```

Spoofing a MAC ID

It is a real hassle for cable Internet customers that they are restricted to using a single machine for Internet access. If you want to plug your laptop in your friend's cable Internet connection, you have to call the service provider to refresh the MAC address.

The MAC address is permanent to the hardware and cannot be changed. Since we operate the hardware via the software abstraction layer, it is quite possible to do some software-level cheating for the network card's MAC ID. We can simply spoof it to some other MAC addresses.

You can obtain the original MAC ID from the *ifconfig* output. Mine is as follows:

```
eth0 Link encap:Ethernet HWaddr 00:1C:23:FB:37:22
```

Now, let's change the last part of the MAC ID from 22 to 23:

```
ifconfig eth0 hw ether 00:1C:23:FB:37:23
```

Now, run *ifconfig* again:

```
[root@gnubox slynux]# ifconfig eth0
eth0 Link encap:Ethernet HWaddr 00:1C:23:FB:37:23
    inet addr:192.168.0.2 Bcast:192.168.0.255 Mask:255.255.255.0
    BROADCAST MULTICAST MTU:1500 Metric:1
    RX packets:0 errors:0 dropped:0 overruns:0 frame:0
    TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
    collisions:0 txqueuelen:1000
    RX bytes:0 (0.0 b) TX bytes:0 (0.0 b)
    Interrupt:17
```

Easy enough? Well, let's consider the following instance: suppose you are in a Wi-Fi campus and the access to the wireless network is restricted using MAC addressing. You can simply look at your friend's laptop MAC ID and spoof it. Yes! You are now free to access the network. Have fun!

DNS (Domain Name Service)

DNS is responsible for name resolution. When you point your browser to *www.google.com*, it points to a server on the Internet. How does that happen? As you are aware, all networked computers are assigned with IP addresses. But how do you access the Web page hosted in Google's remote machine by simply typing a name like *google.com*?

That phenomenon is achieved using domain name resolution. There are some servers on the Web called name servers (or DNS) that resolve certain names to corresponding IP addresses, like *google.com* to its corresponding IP address, in our case. So, we should have the IP addresses of the DNS servers (generally provided by the Internet service provider) handy so that we don't have to remember everyone else's when we browse the Web. When we point our browser to *google.com*, it consults one of these name servers to find out the IP address and thus load the Web page. But where do we configure the IP addresses of these name servers?

If your network is configured with DHCP, there is no need to specify the name server explicitly. For static IPs, it is, however, necessary. We enter the DNS servers' IP addresses in the */etc/resolv.conf* file. Mine looks like the following:

```
nameserver 208.67.222.222
```

```
nameserver 208.67.220.220
```

Note that you don't really need to use the DNS addresses provided by your ISP. For safety purposes, I use OpenDNS—the IP addresses are listed in the above snippet. You can learn more on why OpenDNS is a much safer bet at *www.opendns.org*.

SSH (Secure Shell)

SSH can be defined as the blood of *nix networks. SSH enables users and administrators to make remote logins to other machines that are connected through any kind of network. If you know the user name, password and IP address of another machine on the network, you can remotely log in to that machine and work on it as if you are actually working in front of *that* machine. The following is an example in which I'm authenticating to a system with the IP address of 192.168.1.3 as the user test:

```
[root@gnubox ~]# ssh test@192.168.1.3
The authenticity of host '192.168.1.3 (192.168.1.3)' can't be established.
RSA key fingerprint is 9f:61:ae:ac:8f:75:bb:3a:02:4a:f4:6c:7d:b9:0d:07.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.1.3' (RSA) to the list of known hosts.
test@192.168.1.3's password:
-sh-3.2$ echo I am on 192.168.1.3 Machine
I am on 192.168.1.3 Machine
-sh-3.2$
```

You can open the CD tray of the other machine, close the tray, shutdown, reboot the machine -- depending on the privileges the user name you've logged in with, has.

sftp is an extension to the *ssh* protocol that helps us to use the SSH connection to transfer files between machines. The following is an example:

```
[root@localhost ~]# sftp test@192.168.1.3
Connecting to 192.168.1.3...
test@192.168.1.3's password:
sftp> ls
Desktop  Documents Download Music  Pictures Public  Templates
Videos  a.out  test.bin  file.cpp t.c
sftp> get t.c
Fetching /home/test/t.c to t.c
/home/test/t.c          100% 239  0.2KB/s  00:00
sftp>
```

To download a file from the remote machine we use the *get* command, and to upload a file, we use *put*. In the above snippet you can see that I'm downloading a file named *t.c* using the *get* command, after logging in to the remote machine using *sftp*.

sshfs is another extension to SSH, which empowers you to mount directories on a remote machine as a filesystem to a specified mount point:

```
root@localhost ~]# sshfs test@192.168.1.3:/home/test /mnt/test
test@192.168.1.3's password:
```

In the above snippet, I'm mounting the home directory of the user 'test' on 192.168.1.3 to my local machine under the */mnt/test* directory.

Proxy server configuration

Many of us on a college campus or office network access the Internet through a proxy server. How do you set the proxy server details in your shell environment? You can set the proxy for different protocols as follows:

```
export http_proxy="http://192.168.0.1:3128" ; // HTTP proxy
export ftp_proxy="192.168.0.1:3128" ; //FTP proxy
```

If you want these settings to be permanent, each time you log in add these lines to your *~/bash_profile* file.

That's all, folks! Hope you have enjoyed learning the secrets of networking. Happy hacking till we meet again!




By: Sarath Lakshman

The author is a Hactivist of Free and Open Source Software from Kerala. He loves working on the GNU/Linux environment and contributes to the PiTiVi video editor project. He is also the developer of SLINUX, a distro for newbies. He blogs at *www.sarathlakshman.info*

Are You RHCSS- certified?

A Red Hat Certified Security Specialist (RHCSS) certification ensures impenetrable competencies for the open source architecture.



It too frequently, we are reminded that a networked world brings us into contact with the best and the worst that the world has to offer. Today's organisations must counter the ingenuity and determination of criminals and pranksters with equal ingenuity and determination. The increasingly complex world of IT security requires reliable measures of technical qualifications so that organisations can identify the people who are qualified to implement security solutions.

Relevance of RHCSS to the industry

External security: Unfortunately, the unfriendly world is just getting more

unfriendly, and the stakes are getting higher. Government agencies, particularly those involved in defence, must contend with the threat of cyber-terrorism, sabotage and espionage. Businesses must be on guard for professional intruders who are interested in stealing confidential information.

Internal security: More troubling still, the computers, networks, and Internet access that have made workers more productive have also given the disgruntled or wayward employee new ways to seek revenge and new temptations for misbehaviour. Putting up firewalls to guard against outsiders while leaving internal



networks and systems wide open is simply naïve.

Organisations look increasingly to security certifications to determine who is qualified to truly protect networks and systems. General multiple-choice certifications partially meet the need for such credentials. However, there is a need for a performance-based open source security certification like RHCSS that focuses on technical implementation, rather than policy or theory.

To this end, Red Hat offers the Red Hat Certified Security Specialist (RHCSS), a performance-based security certification that requires advanced skills using Red Hat Enterprise Linux, SELinux, and Red Hat Directory Server. An RHCSS can help your company meet the security requirements of today's enterprise environment.

Red Hat has long attracted industry notice because of live, performance-based testing in certification programs like Red Hat Certified Technician (RHCT), Red Hat Certified Engineer (RHCE), and Red Hat Certified Architect (RHCA). RHCSS joins this family as a first-of-its kind performance-based security certification.

What's expected of a certified RHCSS?

An RHCSS certified person is one who has demonstrated the expertise on technologies used to secure the Red Hat Enterprise Linux (RHEL) operating system. Those certified have demonstrated skills in setting up a centralised authentication and user identity management in a heterogeneous environment using Red Hat Directory Server. As part of the certification process, these specialists will have written Mandatory Access Control policies in SELinux to lock down applications running on RHEL. The RHCSS modules also cover IPA (Identity, Policy, Audit), an integrated security information management solution combining Linux (RHEL), Red Hat Directory Server, MIT Kerberos, NTP and DNS. It consists of a Web interface and command-line administration tools. Currently, it supports identity management with plans to support policy and auditing management.


Why industry would prefer an RHCSS certified person

Industry has grown very rapidly in the past 10 years. Information security is quickly developing into a profession with a strategic focus that has grown beyond its IT roots. Certifications are a measure of experience and knowledge, and businesses increasingly like to see that certification in their employees.

Key aspects from a student's viewpoint

Information security is evolving as a career path. Changes in technologies, business needs, global threats and regulatory regimes have all had an effect on what it means to be an information security professional. One of the most direct paths to professional success in the information security field is through education.

One of the biggest changes we're seeing is that information security is being integrated into part of the overall risk management of an organisation. Overall, this means that information security will mature as a well-defined and strategic career path.

As we define what it means to be an information security professional, we must include the various certifications and understand how they fit into our profession. Being able to certify our capabilities is crucial to our ability to succeed. Whether you gain certifications for technical, management, policy and/or strategic skills, they are all key career building blocks. Whether you are a technical professional or a management/executive professional, certifications must be part of your career planning to put you on the path to success in this unique and challenging profession. 

How to get certified

Current RHCs will obtain RHCSS certification upon successfully passing the following three Expertise Exams:


- EX333: Red Hat Enterprise Security: Network Services
- EX423: Enterprise Directory Services and Authentication
- EX429: SELinux Policy Administration

Examination fee

- EX333 - Rs.10,000/-
- EX423 - Rs. 7,500/-
- EX429 - Rs. 7,500/-

Further details

Visit www.redhat.in/training/certification/rhcsc, or write to training-in@redhat.com. To locate the nearest training institute visit www.redhat.in/training/locations.php



Industry News

Vietnam government opts for FOSS

According to a report by VietnamNet [english.vietnamnet.vn/tech/2009/01/822425], Vietnam's minister of information and communications has issued an instruction for IT workers of government agencies to switch to open source software. As per the report, "...by June 30, 2009, 100 per cent of the clients of the IT divisions of government agencies must be installed with open source software; 100 per cent of staff members at these IT divisions must be trained in the use of these software products and at least 50 per cent use them proficiently."

Some of the software specifically suggested in the news report are OpenOffice.org as the office suite, Thunderbird for e-mailing, Firefox for Web browsing, and a Vietnamese typing software called Unikey.

In addition, the state agencies have been asked to "increase the number of documents and information exchanged among them that are processed by the above software," with a directive that all staff at these agencies must be able to use FOSS at work by December 31, 2010. The minister has even instructed the PC vendors to install open source software, rather than 'cracked' software, on the systems they sell.



Qt to have LGPL option too

Nokia has announced that its Qt cross-platform user interface and application framework for desktop and embedded platforms will be available under the Lesser General Public License (LGPL) version 2.1 from the release of Qt 4.5, scheduled for March 2009. Previously, Qt has been made available to the open source community under the General Public License (GPL).

The move to LGPL will provide open source and commercial developers with more permissive licensing than GPL and so increase flexibility for developers. In addition, Qt source code repositories will be made publicly available and will encourage contributions from desktop and embedded developer communities. With these changes, developers will be able to actively drive the evolution of the Qt framework.

Qt 4.5 will also be available under commercial licensing terms, while licensing for previous versions of Qt remains unchanged. In addition, service offerings for Qt will be expanded to ensure that all Qt development projects can have access to the same levels of support, independent of the selected licence.

"Qt being available under the terms of the LGPL streamlines the licensing of applications built using KDE components on top of Qt-based applications," said Sebastian Kügler, KDE e.V. board member. "This more permissive licensing will further lower the barrier for adoption of Qt and KDE technologies. The KDE team welcomes opening up the development process and is looking forward to improved collaboration between KDE and Qt Software."

For more information visit qtsoftware.com/licensing.

Microsoft works late nights to enable Linux users?

With Barack Obama as the US president, it looks like a change has certainly come. According to a ZDNet.com report [blogs.zdnet.com/open-source/?p=3303] Microsoft apparently worked late night together with Novell's Mono and Moonlight developers on January 19 to enable Linux desktop users to view the US presidential inauguration on Moonlight—the open source alternative to Silverlight.

The US presidential inauguration committee apparently had chosen Silverlight to broadcast the inauguration events over the Internet on January 20. Considering the software only supports Windows and Mac, with the FOSS alternative Moonlight not ready yet to support the live Webcast (as initially admitted by Moonlight chief, Miguel de Icaza), it obviously attracted the flack of the Linux and other FOSS OS users.

However, a surprise came early on the morning of January 20. de Icaza wrote in his blog [tirania.org/blog/archive/2009/Jan-20.html]: "I just wanted to confirm that you can watch today's Barack Obama Official Inauguration video stream using Moonlight on Linux/x86 and Linux/x86-64 systems. All you need to do is to go to the Moonlight Download page... Microsoft worked late last night to get us access to the code that will be used during the inauguration so we could test it with Moonlight."

Looks like the two are in fact working on interoperability. But wouldn't it have been better if pic2009.org had chosen a software/technology that's universally available so none of these issues would have cropped up at the final moments? Or at least, MS and Novell could have collaborated in advance to avoid keeping the FOSS OS users waiting till the last moment.

News Briefs

Sun acquires Q-layer

The Q-layer organisation, which is based in Belgium, will become part of Sun's cloud computing business unit that develops and integrates cloud computing technologies, architectures and services. According to Sun, the Q-layer technology simplifies cloud management and allows users to quickly provision and deploy applications, a key component in Sun's strategy to enable building public and private clouds. For more information, check out sun.com/cloud.

Greg Symon joins Red Hat

Red Hat has named Greg Symon as the vice president and general manager of North American sales. With more than 25 years of business and sales experience, Symon will play a key leadership role in the development and execution of Red Hat's North American sales strategy and growth, the company said.

Symon held various senior sales management and business development positions during a 22-year tenure with the Intel Corporation.

vxVistA to be released under EPL

DSS, Inc. will open the source code for its vxVistA electronic health record (EHR) framework. With this, DSS has effectively removed the greatest obstacle to collaboration in the VistA community by providing the enhanced version of VistA under a commercially-friendly open source licence—Eclipse Public Licence—that can be used to unite the VistA community.

Alan Cox leaves Red Hat for Intel

After 10 years as a kernel developer at Red Hat, Alan Cox has finally bid farewell to the open source vendor. In an e-mail dated December 23, 2008, with a subject line, "Moving on from Red Hat" to the editors of LXR.com, Cox wrote, "I will be



departing Red Hat mid-January having handed in my notice... I'm leaving on good terms and strongly supporting the work Red Hat is doing. I've been at Red Hat for ten years as contractor and employee, and now have an opportunity to get even closer to the low level stuff that interests me most. Barring last minute glitches, I shall be relocating to Intel (logically at least, physically I'm not going anywhere) and still be working on Linux and free software stuff."

Although Cox is leaving on good terms with Red Hat, as indicated in his e-mail: "I know

some people will wonder what it means for Red Hat engineering. Red Hat has a solid, world class, engineering team and my departure will have no effect on their ability to deliver." But his reason for leaving the company — "to get even closer to the low level stuff"—makes us wonder about the shift of focus of the Linux vendor. It seems like they are more focused on the applications stack now. However, this doesn't mean Red Hat is likely to abandon its focus on kernel development any time soon.

SCO announces the auctioning of two business divisions


The SCO Group, the company that sued IBM over copyright infringement of its UNIX code in 2003, has announced that it has filed a formal 'Reorganisation Plan' and disclosure statement with the bankruptcy court in Delaware, the United States, on January 8, 2009.

As part of the proposed reorganisation plan, SCO intends to conduct a public auction to secure additional operating funds and investment in its OpenServer product line as well as its mobile business.

The company believes that the auction will maximise the value of an asset sale and will ensure the ongoing development of these businesses. Several investment groups have expressed interest in acquiring assets via the public auction, claims the company. Jeff Hunsaker, president and chief operating officer, SCO operations, said, "One goal of this approach is to separate the legal defence of our intellectual property rights from our core product business. The auction process is expected to ensure that the future revenue from the OpenServer and mobile businesses stays with those assets and provides an uninterrupted path forward for our customers, products and employees," said Hunsaker.

The plan is subject to the Bankruptcy Court approval. A hearing for the approval of the disclosure statement is scheduled before the Bankruptcy Court on February 25, 2009.





Analyse Your Disk Space Availability

A sneak peak at three handy tools—the command-line based `du` command, KDE's KDirStat and GNOME's Baobab.

*H*ard disks nowadays indeed come with lots of storage space. However, if you continuously add data without cleaning the unnecessary bits and bytes, some day you will surely run out of free space. I compile the *LFYCD*, and therefore need to download and test a lot of software. I always try to build the software from source. After a successful software installation, there is no need for the built binaries in the source directory. If I forget to clean them, after a few days, I run out of free space. So it becomes necessary for me to find a tool that eases my house-cleaning job. The *du* (disk usage) command and two GUI tools—KDirStat for KDE and Baobab for

GNOME—solve my problems.

du tells me about disk usage

Let's start with the *du* command, which summarises the disk usage of each file and directory on a disk. If you run it without any arguments, it displays the disk usage of files and subdirectories in the present working directory. Now, try the following commands instead (note that my present working directory is my desktop):

```
sandeep@linux-i7c8:~/Desktop> du -ch
4.1M  ./My Stuff/Blender 2.46 review/images
4.2M  ./My Stuff/Blender 2.46 review
5.9M  ./My Stuff/Swap/Avidemux
2.4M  ./My Stuff/Swap/SOng
```



```

5.1M  ./My Stuff/Swap/VLC
14M  ./My Stuff/Swap
27M  ./My Stuff
17M  ./PDF Bangalore
82M  .
82M  total

```

The `-c` option here gives the grand total, and the `-h` option presents the size in human-readable form. What if you want this `du` listing of the file sizes in an ascending order? Try the following command:

```

sandeep@linux-i7c8:~/Desktop> du -ch | sort -n
2.4M  ./My Stuff/Swap/SONg
4.1M  ./My Stuff/Blender 2.46 review/images
4.2M  ./My Stuff/Blender 2.46 review
5.1M  ./My Stuff/Swap/VLC
5.9M  ./My Stuff/Swap/Avidemux
14M  ./My Stuff/Swap
17M  ./PDF Bangalore
27M  ./My Stuff
82M  .
82M  total

```

Compare the above output with the previous one—do you notice that the order of the directory listing has changed now? Piping the previous output to the `sort` command helps achieve this with the aid of the `-n` (numerical-sort) switch, which reorders the output based on the file sizes provided by the `du` command.

In this manner you can check the contents of the largest directory and see if there are any unnecessary files there. If so, delete them to free up space.

Another handy command is `df` that you can use to check the disk space usage on all mounted file systems:

```

sandeep@linux-i7c8:~> df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda1        5.4G  2.5G  2.6G  50% /
udev            247M  180K  246M   1% /dev
/dev/sda3        28G   11G   18G  37% /home

```

Adding the `-h` flag to `df` presents the file sizes in a human readable form, viz. in KBs, MBs or GBs instead of 1K block sizes.

All this is pretty nifty, but let's face it: it's often difficult to work with command line tools when the disk space is large and the number of files are more. In this case the GUI tools such as KDirStat for KDE and Baobab for GNOME come to the rescue.

KDirStat for directory statistics

KDirStat is a graphical disk usage utility similar to the `du` command with some additional cleanup features. Most of the Linux distributions have this software included in their software repositories. If it is not

**FOSTERing Linux is The Winner of the
'Most Promising Certified Training Partner'
&
The 1st Insitute in India to produce RHCSSs**



Varad Gupta
India's 1st RHCSS



Rajveer Singh
India's 2nd RHCSS



Ashutosh Talwal
RHCSS



Praveen Kumar
EX 333



Kapil Duggal
EX 333



Vishal Bhatia
EX 423
EX 429

Bring this coupon
&
Get 10% discount
on
any of our trainings



**The Next RHCSS
This could be YOU!!!**

Get Certified on:

- Red Hat Certified Engineer
- RHCS333 - Red Hat Enterprise Security: Network Services
- RH423 - Red Hat Enterprise Directory Services & Authentication
- RHSA429 - Red Hat Enterprise SELinux Policy Administration
- RH142 Linux Troubleshooting Techniques and Tools
- RH184 Red Hat Enterprise Linux Virtualization
- Clustering with Virtualization
- Mail Server
- Bash Shell Scripting

FOSTERing Linux
Free & Open Source Training Education & Research

SCO M-22, Second Floor Old DLF Colony,
Sector -14 Gurgaon 122 001,
Haryana, Tel: +91-124-4268187, 4080880

55, Deepak Building 13, Nehru Place
New Delhi 100 019 Tel: +91-30880047 / 49
Email Id: info@fosteringlinux.com

redhat
TRAINING PARTNER




Figure 1: KDirStat in action




Figure 2: KDirStat settings

installed by default, you can install it using the package manager of your distribution.

After launching it, you will be asked for the directory that you want to scan for disk usage. I have scanned the entire root filesystem (/) and the output is shown in Figure 1. The output contains the directories in a sub-tree structure, usage in both graphical and numerical form, the number of files and sub-directories, and the latest modified time within an entire directory tree.

KDirStat scans the directories on your mounted filesystem and also scans the FTP or Samba directories. It provides some predefined cleanup actions like compressing it to a *tar.bz2* archive, deleting a file or a directory, or opening the directory in a terminal or in Konqueror. You can edit the cleanup actions or define your own set of actions. These settings are shown in Figure 2. If you are a systems administrator, you can even send an e-mail to the owner of the directory that's consuming a lot of space, notifying him to clean up the unused files. The option is available under the *Report* menu.

I'd recommend you play around with this software to understand its various features and functionalities.




Figure 3: Filesystem disk usage scan using Baobab

Baobab from the GNOME userland


If you are a GNOME user, you've something called Baobab (a.k.a Disk Usage Analyser) to analyse the disk usage available by default. You can run it from *System→Administration→Disk Usage Analyser*.

It's similar to KDirStat. It scans the entire file system or a specific directory provided by you (local or remote). The output for the entire file system scan is shown in Figure 3. The output contains the directory tree, size and the total number of objects. It also displays the size for hard links that exist in a particular directory.

Tips to get some free space

If you install software from sources like I do, then to build it you obviously need to uncompress it and then install. After a successful installation, remove either the tar file or the uncompressed folder. Try to find out the old downloaded ISO images, etc, and delete them if they're not useful. Burn the required ISOs to CD/DVD and then also delete them.

If you're in charge of servers, I'm sure you know that a large amount of space is used for log files. Logs are very useful to track problems or debug issues. But most of the time they are unused. Check in your */var/log/* directory to find out the unused log files and compress them using gzip.

Well, that's all for now! But remember, to get good systems performance we need to keep a balance between disk usage and the free space available. Even though you have a hard disk with a high storage capacity, some day you may run out of free space. So remove the unwanted files from the disk and get more free space. **END** 

By: Sandeep Yadav

The author is a part of the LFY CD team and loves to run *./configure && make && make install* every now and then.



Apeejay College of Engineering, Sohna
in association with Indian Society For Technical Education students chapter presents

CONTRIVANCE-09

Innovations galore...

National Techno-Management-cultural Festival
from February 27 to March 1, 2009

Technical Events

1. Paper Presentation - ISTE Conference
2. Technical Poster Session
3. Open Software (Project Display- CSE/IT)
4. EXEMPLAR (The ECE/E&IE/Mechanical Project Display Contest)
5. INNOVATOR-The innovative product design Contest.
6. OVERNITE- The 8 hours nightlong Programming Contest
7. Creative Animation
8. Code Obfuscation
9. Go Go Go....The Online Gaming
10. Get the Bug Out-Debugging contest
11. Pwned !! (The game development Contest)
12. Circuit designing
13. Action Reaction-Laws of Motion
14. Mathematica – The Mathematics Quiz
15. General Quiz

Management Events

1. Business Plan
2. Business Quiz

Contact:

Ayan Tyagi (President ISTE, Students Chapter) 9899966549

Technical Coordinators:

Himanshu Gupta (9999063021)

Prakhar Srivastava (9896440038)

Literary Events

1. Mock Parliament
2. Speechless The Debate
3. My Perspective The Group Discussion
4. Impromptu The Extempore
5. Spelling Maestro

Workshops

1. Digital Film Making
2. Lectures on Network Security by Purdue University USA
3. Workshop on Code Obfuscation
4. Entrepreneurship Workshop.
5. Workshop on FOSS by Linux For You Magazine.

Cultural Events

1. Rock Band
2. One Act Play
3. Bhangra
4. Western Group Dance
5. Duet and Solo Songs

Literary, Cultural coordinators:

Sulabh Vij (9953184007)

Suchet Kaushik (9711976005)

Hospitality:

Mohit Sharma (9996276465), Harsh Pulkit (9896232674)

Venue: Apeejay College of Engineering, Vill. Silani, Sohna, Gurgaon

Web site: aces.apeejay.edu

Email: contrivance2k9@gmail.com

Exclusive Media Partner:



‘We have all that Red Hat has, but our prices will pretty much be half’

Shane Owenby
Director, Linux & Open Source, Asia Pacific,
Oracle Corporation



In this interview, Shane Owenby, director, Linux & Open Source, Asia Pacific, Oracle Corporation, tells us there is nothing like ‘one’ open source community. He also presents what Oracle has, that he believes is missing in Red Hat, which makes the database giant a better alternative for customers. He shares with us Oracle’s contributions to the Linux world, and expresses his surprise at how little credit the company gets for it.

Q What role does open source play in the overall Oracle game?

Well, we are clearly known as a private source company—the word ‘proprietary’ seems too harsh to us. I also think the term ‘open source’ is too big for us. We don’t think that there is just ‘one’ open source community any more. Back in the days when it was very specific, all the open source communities were pretty similar. Now you’ve got very large corporations such as Oracle and IBM that have done some massive things in many different areas. Open source has diversified so much that we cannot refer to something as “the open source community” any more.

We evaluate open source projects based on our contribution to them or their usage. There can be some projects that we adopt; some that we will adopt and integrate; some that we will adopt, integrate and then contribute to; and there will be those for which we will take money and then provide support. So, each individual project has different

levels of commitment from our perspective, and then we evaluate them on a case-by-case basis.

Q How has the feedback been, considering Oracle chose to base its enterprise Linux distribution on RHEL?

Customers look for a single-solution support vendor, and if they get everything from one vendor, it is magic.

The enterprise customers want one version of Linux so that there is hardware and software compatibility. Our approach is to offer enterprise customers the entire stack and not just one component of that stack. At Oracle, what we do is called validate configurations, which includes storage, servers, database applications and, obviously, operation system versions.

When customers come to us and ask for an ERP solution, we have a solution for them. But what kind of an answer does Red Hat have for the same customer? Typically, “I don’t know, but I can tell you this pack of Linux is great for this piece of hardware.”

Q But compared to the experience Red Hat has had with Linux and open source, how do you slot yourself on the same level of expertise?

In the late 90s, there was something called ‘network computers’. That was a long time ago. What people don’t know is they were actually based on Linux. We developed that network computer based on Linux, and that is the Linux engineering team we still have today. Lots of people in the industry don’t know that we were providing patches to Linux for our large Oracle customer base even though they had active Red Hat subscriptions. It was long before we announced our Linux programme.

So, you can say that we have been in the Linux support business unofficially for quite some time. You think Red Hat has got all those big names, like Alan Cox [editor’s note: till recently, as he’s joined Intel in January 2009], contributing to it. That’s true! But we too benefit from the same engineers. And I don’t think that we are at a different contributing level from Red Hat. If you look at the contributors to the Linux Foundation, Oracle is quite high on that list. And if people still say that we don’t contribute that much, it’s not fair.

If you look at the large number of projects in the open source community that are freely running, Berkeley DB is a major one. It is a major contribution we have made—released under a dual licence. Although Oracle has been a good open source corporate citizen, I can’t figure out why people don’t want to give us credit for it.

Q Oracle Enterprise Linux is based on RHEL, so is CentOS. And when people talk about RHEL alternatives, or clones, CentOS comes to mind immediately. Oracle Enterprise Linux is pretty much unheard of.

What does CentOS do? They pick up RH, remove logos and trademarks, and redistribute it. In our view, there should no longer be a CentOS. It takes more time to consolidate servers on the Linux side, and as much as we can, we are doing exactly the same thing. But we are investing a lot of technical resources also.

Q What kind of presence do you have in the LUGs around the world?

What kind of people go to the LUGs? Are the enterprise customers we are talking about, those who are on top of the pyramid, there? In LUGs you go about spreading the word about what you are doing. That is not a core focus of our market. Is it important? Yes, of course it is, if there is an Ubuntu vs Debian vs CentOS sort of debate. But this is not the kind of market we are in. Our sweet spot is the enterprise customer who wants our support.

Q But don’t you think a lot of developers who influence decision makers are active LUG members?

If we had the resources we will also go to the LUGs but, as I said, that is not our focus area. If I have to choose between LUGs and CIOs, I will go the CIO path. It’s not that we are denying their importance, it’s just a matter of

the bandwidth that we have. Most of the LUGs are also non-Oracle shops, generically.

Q So, what do you have that other Linux vendors don’t?
We’ve a feature called premium backporting. So, imagine that this October 1, you have a project to start that is meant to end on December 31. In these three months, you do all kinds of testing. Two months later, after a hard-core testing phase, you find a bug that’s affecting the entire stack. What are the options you’re left with now to fix the problem?

Traditional Linux vendors will say upgrade to the latest version of that package and you will inherit whatever has been changed in that package—enhancements, new features, other bug fixes... whatever. Oracle customers said this is not acceptable. They don’t want any ‘delta’ between fixing their problem and other stuff, because it increases the risk. They want that one bug to be fixed in that specific version of the package that they are using.

What Oracle says is, if you want to have that level of support, we offer it. It’s never been offered ever before in the enterprise Linux space. If someone can do it, it’s good. But it required a significant amount of resources from the engineering point of view. This is what our customers are saying they need from Linux. A little bit more support to promote that adoption. Okay, we do it, so now you can adopt Linux!

Q But even if you look at Red Hat, they are still supporting RHEL 3 and RHEL 4, with a pretty outdated software stack.

I want to give you some specific examples. Let’s say you’ve got GCC version 2.3, and now you have found a bug that needs to be fixed. But by this time, Red Hat may have added 2.4, 2.5, 2.6, etc. So when you approach them for the fix, their answer will be to upgrade to the latest version of that package for that release. The situation is like, you are at GCC 2.3, the latest RHEL package for GCC is 2.76, or whatever, which means there is a gap there. For an enterprise, every new feature inside a version is a risk. So what you want is that one bug fixed, and not have to upgrade to that latest version packed in the RHEL release. I’m not saying from RHEL 3 to RHEL 5; I’m talking about packages inside those release chains. So wherever we jump from one version to another, there can be several changes. This is a risk. Customers don’t want to see that, but that’s the answer they get from Red Hat.

But what we do is find the bug in that package and fix it in the same version, and give it to the customer. It has never been done by any one, as it involves significant amounts of engineering. If the customer is okay to upgrade to whatever package in that release, then it’s okay; you don’t need that premium backporting feature.

I’m just telling you the difference between Red Hat and us; and there is another very important thing—we offer life-time support. So for concepts like RHEL 3 reaching ‘end of life’, there is no such thing with Oracle Enterprise Linux—we support that forever. The same is the case with

our database. It is company policy. I don't think there is any other Linux distribution that does that.

Another thing is that we let our customers manage their infrastructure for free. Red Hat charges on Red Hat Network Satellite and also has a pre-server fee.

Q This premier backporting feature comes at an extra price tag or is it a part of the regular pricing?

This feature is charged for. Generically, we have a price for one and two CPU machines, and those greater than two. We have two levels—one with lifetime premium backporting support, and the other without it. We also have one for only updates. Here, you can't call us. Now, Red Hat has service 5 days a week, during working hours (12/5, as against 24/7), which we don't have. We think that if you are going to run a production-level server, you should always have 24/7 support. How can you say, "Oh! the bug is not big enough, you can only call us during those hours." From our perspective, that's not acceptable. So we have 24/7 offerings only. All prices are on the Web so you can check it for yourself. We have all that they have, but our prices will pretty much be half of theirs. There're no 'apples and oranges' comparison here. We have premium backporting and lifetime support that they don't. We have 24/7 support, and they don't.

Q What kind of support structure is there for non-Oracle customers who are only buying OEL and not typical Oracle products?

We think you are still an Oracle customer because you are buying Oracle Linux. We don't differentiate among our customers; the process is the same. However, the interesting thing about Linux is that our Oracle support organisation has a deep knowledge of databases, middleware and all that. So, if you identify a filesystem bug, who do you call? Do you call the operating system vendor or the database vendor -- because it's a perfect line between those two?

The database is messing with the filesystems and the operating system kind of controls the filesystem. The Linux vendors don't know about databases, and the database vendor doesn't know about Linux. It's hard to find that level of skills in this area. When you call Oracle, we have support engineers with access to the database, along with others with a Linux background. Our Linux guy says, "Alright, if you're doing this in the database, it can trigger those filesystem issues." So, we are able to answer those questions much faster, just because of the availability of the knowledge inside Oracle. Therefore, the value we offer a customer is: "Don't worry, from the database to all the way down, you call us; it's our headache to sort out the problem." There is extreme value in that.

Q In these times of an economic meltdown, do you see more migration towards Linux and open source?

Yes, exactly! You are paying the same, but the expectations are much higher. So, the business is getting squeezed. How to deal with that? The solution is to bring the same value, but at a much lower price point, and that's where the strength of

Linux comes in. If you look at the financial industry, they were going for big-boxes because that's how they look to scale out as they had plenty of money. Well, the situation is different now, and we see a lot more proof of concepts and a lot more deployment projects. There are a lot of small projects that are precursors to much larger projects. And, we see a significant number of those happening now; of course, the existing Oracle stack of solutions is still there.

Q For your team, is it the market you're after or the development community? For example, there is a team in IBM that taps the development community, and there is another team looking at the business side of it.

Oracle has a team focused on developers. The goal of my team, on the other hand, is to sort of blanket the world with Linux. So, my focus is on the business side. However, in this specific field, developers also play an influential role. They influence the decision makers.


And obviously there is Oracle's strong commitment and history. If you want to go with other operating systems like Solaris or AIX, we still offer support as Oracle. But this is the strategic direction we are going in and I must add that some of our products get released for Linux before any other systems. Our developers internally develop on Linux, so there's quite a lot of momentum there. There is something called Oracle on Demand, where Oracle runs an Oracle customer's Oracle software for them—we run that on Linux. So, we trust Linux that much.

Q What kind of challenges do you see in people migrating to Linux?

The human brain is amazing—in the future, people always do what they have done in the past; not everything, but a significant number of things, nonetheless. If they have always bought large UNIX boxes, they will think, "Of course, I will buy UNIX boxes."

Now, the economic downturn is helping us. People are now not going for those expensive boxes. They are looking for other ways. And clearly we are the leaders in the other way, because Oracle has been doing clusters and bridges for a very long time. Therefore, combing the human brain with some flexibilities, in addition to the economic downturn are creating opportunities for us.

Q What if someone is not interested in Oracle? Say they are using other databases like MySQL or PostgreSQL, but they want to go with OEL? Do you have the expertise in handling the issues of such a customer?

Sure, we support whatever ships with the RHEL Advanced platform, which even includes GFS (Global File System). We have our own OCFS (Oracle Cluster File System). However, whatever they ship, we support. We are a database company, so we can obviously support a database query.  **END**

An interview by the LFY Bureau folks.



National Conference on Open Source Software

May 25th - 26th, 2009, Mumbai, India

Organised by

C-DAC, Mumbai

Supported by

IEEE Computer Society, Mumbai and Chennai chapters
Computer Society of India, Div II on Software & SIG-OSS

Call For Papers

NCOSS—2009 is a forum to bring together the various groups working on developing Open Source Applications catering to specific domains in the ICT world—education, health, accessibility, localisation, e-commerce, disaster management, expert systems, machine learning, etc. A number of high-quality software solutions are available in many of these areas, for example, SugarCRM, Koha, Drupal, Moodle, Sahana, CollabCAD, etc. Work on these systems require a combination of domain knowledge and development expertise. Much of the public awareness in open source is focussed on desktop, operating system and general productivity tools. With this background, NCOSS-09 has chosen to focus on the layer above this, bringing together groups working on various application domains.

The conference will present experiences in deploying FOSS applications, comparative studies among competing software solutions, efforts in adapting and localising FOSS applications, development of new applications, etc. The conference will consist of the following:

- Invited talks by experts from India and abroad
- Presentation of contributed papers selected based on refereeing by a panel of referees
- Exhibition by industry and academia
- Pre conference Tutorials (on May 24th)
- Panel discussion

TOPICS

Papers are invited on the topics listed below: (Other application areas may also be considered).

Accessibility	Machine Learning and Data Mining	e-Governance
Indian Language Computing	e-Health	e-Commerce
Localisation	Knowledge Management	Disaster Management
e-Learning	Collaboration Technologies	Content Management
Information Extraction and Retrieval		

INSTRUCTIONS

- Papers must report original work carried out by the authors. The work can include enhancing existing Open Source applications for specific requirements, development of new solutions and comparative analysis of competing solutions. Direct survey or overview papers are not acceptable.
- Length should not exceed 10 pages of A4 size in length (approx. 5000 words) including figures, etc.
- Papers should be in English.
- An abstract of about 100-200 words and the area(s) under which the paper can be categorized, must be included with the paper.
- The author names and affiliations along with the main area of the paper should be given only on a separate cover sheet. Papers should be in one of the following formats: PDF, RTF or ODT. Accepted papers will be published in the conference proceedings.

Visit <http://ncoss.cdacmumbai.in> for more details and paper submissions.

Media Partner:





GeeP

The Geeks of Pune!

A forum for Linux experts to share their wisdom, and more...

It was a typical evening at one of the coffee houses in Pune. Most of the tables were occupied by youngsters, with the typical topics of discussion varying from college *katta* news to the latest movies, cricket and politics.

But at one corner there were around 20 geeks discussing hardcore Linux issues over coffee. There was no age distinction in the group and participants ranged from graduate students to industry veterans. The GeeP meet was in progress.

As the name suggests, this group consists

of Linux geeks. Unlike the regular LUGs, 'GeeP' is primarily focused on Linux systems software, which typically includes the Linux kernel, device drivers, GCC and compilers, and other aspects of the system.

GeeP started off in July 2006 as a mailing list at tech.groups.yahoo.com/group/geep_linux. Amit Kale, CEO and founder of LinSysSoft Technologies was the driving force behind the formation of GeeP and making it popular across Pune. Although it started off to get together the Linux kernel developers from around the city, the mailing list grew rapidly and has now more than 177 members,



including folks from other cities like Bangalore, Allahabad, Chennai, etc.

GeeP is a non-profit group. The members have a philosophy of sharing wisdom amongst themselves, while job postings and other hoax stuff on the group mailing list are discouraged.

Typically, group members get together once a month at a Barista outlet. The discussions mostly tend to focus on Linux systems issues, with members sharing their experiences and thought processes with respect to their professional life. Of course, members also share other industry news, giving rise to some spicy discussions.

The group plans sessions for its members once a month. In November, it was filesystems, and virtualisation the month before. These sessions are absolutely free of cost and one can opt in just by dropping a mail.

The sessions are typically divided into two parts—the first covers the ‘basics’ for newbies, and the second half is generally an advanced session on the same topic. A GeeP member generally covers the first half, and experts are invited for the second. The presentations of past sessions are available to members through the ‘files’ section of the group. These ‘sessions’ sometimes are hands-on, and attending them can be a real learning experience.

The topics for these ‘sessions’

are collectively decided upon by the group members.

November was really a quantum leap for the group when it formally went online through www.geeksofpune.org. Kale and the rest of the ‘core’ group of GeeP members are optimistic about the future. To quote him: “The group is one-of-a-kind in India and stands out distinctly from regular LUGs. The group seeks affiliations from the industry and educational institutes. We expect a lot of action in the following months.”

GeeP has also started going to engineering colleges, where members introduce graduate students to the use of Linux. Companies that have a significant presence in the Linux/OSS workspace are also in touch with GeeP and often offer their conference rooms for the group’s monthly sessions.

So all you geeks, welcome onboard! 

Links and resources

- Official website: geeksofpune.org
- Yahoo Group: tech.groups.yahoo.com/group/geep_linux

By: Nilesh Govande

The author is a Linux enthusiast and can be contacted at nileshgovande@yahoo.com. His areas of interest include Linux systems software, application development and virtualisation. He is currently working with the LSI Research & Development Centre, Pune.

p00r JOB MARKET ?
Campus Placement ?

Future is OPEN

Now time to
Upgrade your Skills, Register...



RHCE

Only at 

(Red Hat Certified Engineer -
No 1 - Certification in the World)



First Time in S INDIA
Only at 

Zend Certification

ZCE - International Certification for PHP
from Zend Technologies - USA

LAMP Pro

Only at 

(Linux + Apache+MySQL+PHP)
Professional



LPI

First Time in INDIA
Only at 

(Linux Professional Institute - Canada,
International Certification for Linux)

Feb 09 - Schedule

RHCE - 5, 14, 18, 25

LAMP Pro - 7, 11, 18

ZEND Training - 11

Shell & Perl - 7

RHCE Exam Date : 24 Feb 09




Future is open

Lynus Academy Pvt Ltd

Chennai - 044-25021630,
42171278,9952966527

Pondicherry - 0413-2276233,
4202121, 9952966528

contactus@lynusacademy.com
www.lynusacademy.com



Why buy expensive 'branded' servers? Build one yourself using off-the-shelf hardware.

When enterprises want servers, they buy them. But when hobbyists or small businesses like travel agencies and franchises want servers, they can't. All the servers available commercially are blades, razors, racks, cupboards, etc. The fact is: servers don't have to be in huge air-conditioned rooms. They can run on commodity hardware also. We are going to prove that building a commodity server from scratch (you can make and sell these, if you like) using off-the shelf hardware can be done for less than Rs 45,000.

**As Barack Obama famously said:
"Yes, We Can!"**

Yes We Can! With 45 grand, we are going to build a server that supports Web, e-mail, virtualisation, media streaming, DNS, PXE, fax, a file server and VNC server roles.

Because of the budget we have, we are going to assemble our own machine (or at least buy the parts ourselves, and get it put together by the local techie). I will recommend specific hardware, including their model numbers.

So let's start buying:

- Processor: Intel Core2 Quad Q6600 2.4GHz @ Rs 9000
- Motherboard: Gigabyte EG45M-DS2H @ Rs 8000
- RAM: Transcend JetRam DDR2-800 800MHz—2x 2GB modules @ Rs 5500 (Rs 2,750 each)
- Hard Disk: 4x Seagate Barracuda 7200RPM 500GB @ Rs 14,000 (Rs 3,500 each)
- Cabinet: iBall i701 Server Cabinet @ Rs 4,000 (Pricey!)
- Keyboard and mouse: Microsoft

wireless desktop combo 1000 @ Rs 1,850

- Monitor: You can get a 17-inch CRT from Samsung at Rs 3,500
 - Optical: LG GH22NP20 DVD-RAM Drive @ Rs 1,250
- Total cost: Rs 43,600 (tentatively). Prices vary, but there are chances you can get a better deal from the vendors in your city if you have time to find out the prices yourself after visiting a few shops.

But let me tell you, you are not bound to this configuration. At a bare minimum, a 500MHz PIII with 256 MB RAM and 80 GB of hard drive space will do, but you will not be able to run many applications. And if you are a hobbyist, you can get away with VMWare server (not VirtualBox or QEMU because of the performance-hit involved).

So what do we have? 2 TB of HDD space, 4 GB RAM, a gaming motherboard (I chose that one because it can be overclocked if need be), a 2.4 GHz QC processor that processes four threads per clock cycle, and a cabinet with a very robust SMPS and plenty of drive bays.

You have a very sleek wireless mouse and keyboard combo, but that's something you can cut down on. Because, at the end of the day, the server will have no monitor and input devices -- everything will be managed via remote X.

Oh, and of course, the UPS system is extra!

The ribcage

To build the basic backbone, or by our analogy, the ribcage of the server, we are going to use (you guessed it!) GNU/Linux. For our distro of choice, we are using a desktop distribution, Kubuntu. You can use Xubuntu or plain and simple Ubuntu as well.

Because this is a very complex procedure, we are going to do this in stages.

Stage 1: Installing the OS

First, get the latest CD of the distro you are going to use. Head to ubuntu.com, kubuntu.com, or xubuntu.org and get a CD. And in case you thought you are going to use the *LFY* December 2008 DVD, it has an installation bug that affects Xubuntu. Get the desktop edition, not the server edition.

Now switch on your PC for the first time, and set the date and time on your BIOS. Search for something like "Boot device priority", and set your PC to try booting from the CD first, and then the hard drive. Save your settings, and exit.

When you reboot, your PC will get stuck and complain about the absence of an operating system. So, insert the Kubuntu CD and hit the reset button on the cabinet.

Directly go to *Install Kubuntu* and start the installation.

The installation process of Kubuntu is just seven screenfuls long. The process is very simple. The only

The Star Bus Topology

A topology is the arrangement in which computers on a LAN are physically connected together. There are many such topologies, such as Ring, Star, Bus, Mesh, etc. Since the Ethernet is not connection-centric, any sort of arrangement works, but these topologies have a wide use, as they are the only way to keep the physical connections organised.

The Star-Bus topology is a hybrid topology, using the best of both the Star and Bus worlds. In this arrangement, a single high-capacity line, called the Bus, runs from the server to a hub. From the hub, many more lines run to one client each, or to other hubs. In the latter case, the hub-to-hub line becomes another bus. In a schematic diagram, each of these clients on a hub is represented by spokes running from the hub to the client; thus the Star.

In a Star-Bus topology, you have flexibility as well, because you can have different servers at different levels, each tending to the levels below it.

place where you need help is for partitioning.

Create a new MBR on */dev/sda*, and create a 500 MB *ext2* (not 3) for */boot*, a 40 GB XFS for */* (root) and a 6 GB swap space. After that, devote the rest of */dev/sda* to a partition mounting under */tmp/ServerTemp*. Use XFS here and then the manual partitioning option.

Complete the installation as usual and reboot into KDE4.

At this point, your server should not be connected or configured for any sort of networking at all.

You will be working as the root user throughout. This is a big security issue, I know, but this is the only hassle-free way of getting everything done perfectly.

Oh, and make a note of disabling all power saving schemes and screensavers. Since this is a server, it may need to stay up for years at a stretch, often without months of desktop activity.

Stage 2: Enable root

Enabling the root isn't a big issue. Open Konsole and type in *sudo passwd root*. Enter a root password twice. Bingo! The root is enabled now.

Stage 3: Internet connectivity

The time has come to connect our server to the Internet. I assume that you are using an 'unlimited' usage broadband connection that has a static IP address. If you don't have a static IP address, ask your ISP for one and write it down somewhere. Having a static IP address is very important.

Now ask your ISP service person to configure your modem to be always-on, i.e., you don't need to enter a password to access the Internet. This constitutes entering the user name and password into the modem itself so that it authenticates every time it switches on.

Run a length of RJ-45 wire from the port on your modem to the network card socket on your motherboard.

Your NIC will have a device node `eth0`. Add the following two lines to the `/etc/resolv.conf` file:

```
nameserver 208.67.222.222
nameserver 208.67.220.220
```

Execute the following command:

```
ifconfig eth0 192.168.1.2
route add default gw 192.168.1.1 ## assuming this is your
gateway address—the IP address of your router
ifconfig eth0 up
```

Now reboot the system and try to browse the Internet. Everything should work.

At this stage, run an `apt-get update` and `apt-get dist-upgrade` to upgrade all installed packages to their latest versions.

Stage 4: Let the world welcome you

Yeah, you got it—DNS.

You can have your own DNS if you want, and I'll cover that in the last part of the series. But right now, we will not have our own DNS server. What we are going to do is sign up for a DNS service provider.

First, register a domain. You can get one for free at www.co.cc. Sign up for a free account. You can have up to two domains per account. They are in the form www.yourdomain.co.cc. Mine is www.theossfreak.co.cc, but don't freak out here. Not all domains are free, some require up to \$900 per year. The trick is to make a simple spelling mistake. MangoesAreSticky comes for \$7 a year, but take the e out of Mangoes and the domain is free for life.

When you set up your domain, use the first of three configuration options. Set up the nameservers as `ns1.everydns.com` through `ns4.everydns.com` (4 in all, you need to add 2).

Now go to www.everydns.com. Sign up for another free account. You will be signed in immediately. Add your domain as a basic domain (on the left). Give the full name without the `www`, like yourdomain.co.cc. The right pane will now show a table with a single 'A' (the IP address) and a 'CNAME' record. You need to add one more first and then remove the other two.

Remember the static IP address that your ISP gave you when you asked for it? Just below the table is the form for adding a record. Give the domain name without the `www` in the first field. Now let the record type be A. In the value field, give your static IP address. If it's BSNL, then it most probably starts with 59. Then commit the changes. Your existing A record will change to reflect the new settings with a TTL or Time To Live of 3600 seconds, or 60 minutes. We are done.

What now?

Because server connectivity is brilliantly tough, you are bound to hit problems. You can try out Internet forums, try your own luck, or you can e-mail me. The problems vary greatly and are dependent on many things, such as the USB drive not being found, or you not being able to connect to the Internet. Take these problems head-on and overcome them. Remember that this server has no enterprise support.

We are not building this server as an expensive Internet browsing machine. We will eventually set up the server role-by-role. I did want to set up every piece of software from the sources, but ran out of patience. Hundreds of dependencies, and equally high compile times, made it an exhausting task. Thus the controversial choice of the OS. (Whod use a desktop OS as a server?)


There are advantages to using Ubuntu. It's a very simple to use Debian-based distribution. We can tweak it exactly to our needs. Even now if you think KDE4 is too heavy for a server, issue `apt-get install xubuntu-desktop` to go with XFCE 4.4.2. Brilliant!

Next month, we will start setting up our server. We will first set up the DHCP services. Then comes the setting up of a TFTP server as the first real server role. (TFTP is kind of useless, but it is used to remote-boot PCs over the network using PXE.) We will also set up a firewall, following which we will set up a basic LAMP (Linux-LightTPD-MySQL-PHP-Python3-Perl) stack for our Web apps. (If you doubt LightTPD's power, it is used by both Wikipedia and YouTube, and they're seldom, if ever, down!)

In the subsequent months, we will start by setting up a complete Web and database server, an e-mail server, fax, IRC and Jabber (for chat), virtualisation, file storage (FTP), print and media streaming, in that order. You can skip any function you like, and you can ask me to add roles that you may want. I'm open to suggestions and contributions, in the true spirit of FOSS.

For now, go and purchase another network card, and a network hub. Configure the new NIC as `eth1` with an IP address of 192.168.1/24. This will be the IP address that all the PCs in your network will know as the server. Connect `eth1` to the hub. All the PCs in your network will connect to the hub in a Star-Bus Topology [see sidebar].

I also suggest that you do some research on basic networking, the Ring and Star topologies, and CSA (Client Server Architecture). This will help, as you will know what you are doing.

See you next month!  **END**

By: Boudhayan Gupta

The author is a 14-year-old student studying in Class 8. He is a logician (as opposed to magician), a great supporter of Free Software and loves hacking Linux. Other than that, he is an experienced programmer in BASIC and can also program in C++, Python and Assembly (NASM Syntax).

Delhi University Computer Science Society

in Association with



Presents

Sankalan-09

Compiling Innovations...

- Sprint: C/C++ Programming
- Java Juggling
- Alcoholics
- Techno Speak (Paper Presentation)[†]
- Mind Your Language: Debugging
- Spin a Web
- Open Your Mind: On OpenSource Technology
- Who Dare Wins
- Select * from Brain
- Mind Matters
- LAN Gaming
- Just A Minute
- Turn Court

On 7th & 8th February, 2009

Eligibility :

**B. Sc./B. Tech./BCA/
M. Sc./M. Tech./MCA**

Venue:

**Conference Centre, North Campus,
University of Delhi, Delhi-110007**

For Details Contact :-

Rohit Ahlawat (President)
Mob: 09868774279

Nikhil Malik (Secretary)
Mob: 09990903913

Trasha Gupta (Treasurer)
Mob: 09868791924

Website : <http://cs.du.ac.in/sankalan2009>
E-mail: sankalan2009@gmail.com

Key Sponsors



Co-Sponsor



Magazine Media Partners:



[†] Participants are expected to present an update on state - of - the - art on any topic in IT domain
^{**} Sponsorships still open



Metalinks for Download Anyone?

Still stuck to http/ftp downloads or have you moved to torrents instead? What if I say you can have the best of both worlds?

Downloading large files from a server with http/ftp is becoming obsolete now. People are more inclined towards using torrent downloads because of the obvious advantages over server downloads such as speed limits, connection timeout errors, etc.

Torrents, however, are not so useful when downloading something new like, say, a new Linux distribution release. The seed:peer ratio is almost zero. The servers are overloaded. You neither get much speed from any particular server nor from torrents. How

many times have you wished that you could supply more than a single server to your download manager to increase the speed a bit? Maybe mix torrents with the server downloads for even higher speeds? This is where Metalink steps in. Anthony Bryan came up with Metalink 11 years ago, but it still has not received the limelight it deserves.

In the simplest words, a Metalink is an XML file. This file contains the links to the various mirrors that host the file under consideration. It might also contain the location of the torrent of that file and/or

magnet links. It also has an added advantage when used with some clients. The file can be simultaneously downloaded from various locations in the form of segments (like hash pieces in a torrent, but not the same). It also features automatic checksum verification of the downloaded data and, if any errors are found, there's an automatic error correction feature also. Metalink supports both MD5SUM and SHA1SUM checksums as well as PGP signatures, which are embedded in the *.metalink* file itself.

A checksum is like a signature or a fingerprint of the file. If the downloaded file has no error then the checksum in the Metalink file will match the checksum of the downloaded file. In case the matching fails, the Metalink download client only downloads the segments with errors using different tools like *rsync*. For some clients, you may have to manually download the missing/corrupt data using tools such as *rsync*.

Creating Metalinks

Most distributions these days also provide a Metalink to the ISO files. If you do not have the Metalink or cannot find one, you can create it yourself. There are many Metalink generators available online or you can install one if you want. They generate the Metalink XML file. It contains XML data with a *.metalink* file extension.

The online Metalink generator at the official site of Metalink is good enough for generation and can be found at www.metalinker.org/generator. If you would like to install a generator or use local metalink generator scripts, then you can find those at www.metalinker.org/implementation.html.

For the more geeky folk, or those who want to have complete control over the Metalink file (there is not much to have control over, but just in case), the general Metalink XML format is given below:

```
<?xml version="1.0" encoding="UTF-8"?>
<metalink version="3.0" xmlns="http://www.metalinker.org">
  <files>
    <file name="example.ext">
      <verification>
        <hash type="md5">example-md5-hash</hash>
        <hash type="sha1">example-sha1-hash</hash>
        <signature type="pgp"/>
      </verification>
      <resources>
        <url type="ftp" location="us" preference="90">ftp://ftp.example.com/example.ext</url>
        <url type="ftp" location="uk" preference="75">ftp://ftp.example.net/example.ext</url>
        <url type="http" location="us" preference="80">http://example.com/example.ext</url>
        <url type="http" location="de" preference="10">http://example.net/example.ext</url>
        <url type="bittorrent" preference="100">http://example.org/example.ext.torrent</url>
      </resources>
    </file>
  </files>
</metalink>
```

```
<url type="rsync"/>
<url type="magnet"/>
<url type="ed2k"/>
</resources>
</file>
</files>
</metalink>
```

This is just the general format. You can also add additional fields like OS (OS specific files...the download manager will check your OS and automatically download the file tagged with your OS). The preference parameter has a value from 1 to 100. If a segment is available at multiple locations, then the segment will be downloaded from the location with the maximum priority. There is a provision to add publisher information like name, website, etc.


Downloading files using Metalinks

You cannot use your regular download clients to download files using Metalinks unless your client supports the Metalink format. If you are using Linux, then I would suggest using the *aria2* [aria2.sourceforge.net] download client. It has a command-line interface and comes packed with many features. (Almost all features supported by Metalink are also supported by *aria2*.) The DownThemAll Firefox extension can also handle Metalinks and is my personal favourite [www.downthemall.net]. If you would like to try a different client, the list of all clients supporting Metalink can be found at www.metalinker.org/implementation.html.

For the win formula

Metalink is a cool feature for downloading large files. The ability to avoid connection timeout errors and the package integrity check is something that gives Metalink some added bonus points. Plus, by grabbing chunks of files from multiple sources at once, it makes efficient use of available bandwidth.

Metalink can be better, of course. It still lacks some important features such as provision of different versions of the same file and an option to select one you want. One more feature lacking is dynamic preference value generation, depending on the load on the servers in the Metalink file. The server with the least load must be given first preference.

The advantages still outweigh the disadvantages, by far. Metalink is, as I said, a cool feature. I am amazed how I failed to notice it for so long.  **END**

By: Aditya Shevade

The author is a National Talent Scholar and a third year electronics engineering student, who takes a keen interest in programming and electronic design. A Linux user for the past two years, he enjoys playing the keyboard and is a good photographer. To know more about him, visit www.adityashevade.com.



Niyam Bhushan

If You Love Someone, Set Them Free

If you are her man, give her a mouse.

She moves in mysterious ways. There she is. Online. Giggling to a funny video on YouTube, gossiping on Facebook, and as usual, addicted to online chat and e-mail. She's the one, cajoling him to surf for romantic holidays and book travel tickets. She's one in a million. Actually, one in around 12.32 million, according to a study by the Internet and Mobile Association of India [www.iami.in]. I expect those figures to be significantly higher, since the study is about three years old. Some more recent figures, published at internetworldstats.com, offer different figures: a total of 60 million Internet users in India, which is just 5.2 per cent of India's population. Of these 60 million people, 85 per cent are males, as gleaned from several other sites. Hence the mystery. Only one thing is certain: there is a surge in the number of women embracing the Internet, even though the overall percentage is still low.

Ping for pink

In my personal observations, only a small percentage of women in India own a personal computer. They often share someone else's computer at home, or use a cyber cafe. But if you love someone, you could set them free. Really surprise her this Valentine, by gifting her a personal netbook. These laptops are small enough to fit into a lady's purse, and come in attractive colours. The good ones come pre-installed with GNU/Linux, and if you're lucky, you may just find a majority of models between Rs 9,000 to Rs 21,000. I expect these prices to further fall down by a few thousand rupees through this year. Netbooks suit her lifestyle, too. These can be used anywhere, are more robust than regular laptops, and have less moving parts. She can use them to share photos from a connected digital camera; browse her music on her digital music player; for a little bit of college or office work; and of course, for all her socialising and shopping on the Net.

Surprisingly, the response to netbooks that I've noted among women is quite different and far more positive than those of men. You'd expect them to shirk away from pre-installed GNU/Linux. On the contrary, they just love the idea of not having to deal with viruses and crashes any more, and appreciate how its adoption

drives the price down further for them. Try it. Take her to any of the new gadget shops dotting urban Indian malls, and let her play with one of these petite laptops.


Share the happiness

Don't despair if you can't gift her a netbook in this global meltdown. Buy her a USB thumbdrive with the highest capacity you can afford beyond 2 GB. These may cost you between Rs 400 to Rs 2,000, about the same as a date at a cafe, or a candlelight dinner at your favourite restaurant. Then head over to pendrivelinux.com, and install a bootable GNU/Linux system on the device. Load it up with all her personal data files, such as her favourite music, personal photos and movies. If she's smart, she might just study the smile on your face and ask you: "Is that a USB thumbdrive in your pocket, or are you just happy to see me?"

That's your cue to share your happiness. Show her how she can just plug in her new USB thumbdrive to any desktop or laptop computer and boot into her own

personal GNU/Linux system. She has the freedom to do everything that she always wanted to do with her own personal laptop or desktop, without touching the data on the host computer's hard-disk. She can also save all her data

and files on her thumbdrive, so no one can snoop into her personal life, or accidentally delete her files.

When you really love someone deeply, the most precious thing you can share with each other is freedom. That is when you really discover the difference between true love and dependency-hell. Share with her the taste of freedom: give her the gift of GNU/Linux. For women have been the custodians of free knowledge and wisdom, in an unbroken chain lasting thousands of years from generation-to-generation. Happy Valentine's Day! **END** 

"Is that a USB-thumbdrive in your pocket, or are you just happy to see me?"

About the author:

Inspired by the vision of Osho. Copyright February 2009: **Niyam Bhushan**. freedomyugs at gmail dotcom. First published in LinuxForYou magazine. Verbatim copying, publishing and distribution of this article is encouraged in any language and medium, so long as this copyright notice is preserved. In Hindi, 'muft' means 'free-of-cost', and 'mukt' means 'with freedom.'



**The 3rd World Mobile Summit & Sourcing Fair
The 5th China Int'l Mobile Sci-Tech Exhibition**

Date: 6-7th March, 2009

Venue: Shenzhen Convention & Exhibition Center, China



**Build mobile phone industry brand exhibition,
establish the international communication platform.**



Organizer

China Mobile Communication Association
Shenzhen Municipal Government

Executive Organizer

Shenzhen Zhongrui International Exhibition
Design Co.,Ltd.

Co-hosted

Hongkong Zhongrui Conference Exhibition Limited
Shenzhen Mobile Communications Association

Add: Room 629, West Tower, Tian'an High-Tech Plaza Phase II, Futian District, Shenzhen, China		
Tel: +86-755-82507486	Fax: +86-755-83297552	E-mail: 70211@163.com
Contact: Rebecca Lee	Mobile: +86-13242439028	www.mobilechinaexpo.cn

Indian exhibitors and delegates can contact EFY team for details: nisha.anand@efyindia.com
Tel:0091-11-26810601 Fax:0091-11-26817563

Programming in Python for Friends and Relations—Part 10

Python on the Net Using CGI and WSGI



Let's build a small application of classifieds.

*T*he Internet is wonderful. You can host your application anywhere and access it anywhere! And there is no easier way to share it with friends. It isn't hard to write your own little application. The Apache Web server does all the hard

networking work. Your application will print the output and it will be sent back to the browser that made the request. The concept behind all this is the common gateway interface, or CGI, for short.

Apache in Fedora and Ubuntu is configured to handle Python CGI code out-

of-the-box—only the location where we place the code is different. Let's work with the default locations because configuring Apache is outside the scope of this article.

In Fedora, you write your code in `/var/www/cgi-bin`, and in Ubuntu you write it in `/usr/lib/cgi-bin`. The configuration files are in `/etc/httpd` in Fedora and in `/etc/apache2` in Ubuntu. Try the following code in *hello.py*:

```
#!/usr/bin/python
# Tell the browser that it is just plain text
# \n => a blank line. Indicates end of header
print 'Content-Type: text/plain\n'
# The content
print 'Hello, Friends!'
```

The first line instructs the shell to run this script using Python; but the file *hello.py* should be executable, that is:

```
chmod +x hello.py
```

Now, point your browser to `localhost/cgi-bin/hello.py` and you should see 'Hello, Friends!'.

Getting started with HTML

The Web pages are normally not plain text. They are HTML documents where the tags instruct the browser how to display the content. So, modify *hello.py* to be an HTML document.

```
#!/usr/bin/python
# Tell the browser that it is an html document
print 'Content-Type: text/html\n'
# Print the html document using triple quote for convenience.
print "<html>
<header></header>
<body>
<h3>Hello, Friends</h3>
<p>Welcome - this is a paragraph.</p>
</body>
</html>"
```

The structure of the HTML document is `<tag>content</tag>`, where the content may include other tags. You are concerned with the body part. You now have two statements being displayed—one as a header and the other as a paragraph. You can try different types of headers—h1 through h4 and see the difference.

Displaying a form

You can now build a small application of classifieds for your friends. You will have three fields, a description and a contact e-mail ID. Write the following in *cgi-bin/form.py*:

```
#!/usr/bin/python
```

```
# Define the form
content_header = "Content-Type: text/html\n"
html_page = ""
<html>
<head> </head>
<body>
<h1>Entry Form</h1>
<form method="post" action="form.py">
    <p><label>Title: </label><input type="text" name="title"
value="% (title)s" / ></p>
    <p>Description: </p>
    <textarea name = "desc" rows=4 cols=60 >% (desc)s</textarea>
    <p><label>Email-id: <input type="text" name="email" value="% (email)s"
/></label></p>
    <button type="submit">Submit</button>
    <h4>% (message)s</h4>
</form>
</body></html>
"""

# Python Code
message = email = title = desc = ""
# present the form
print content_header
print html_page % {"title":title, "email":email, "desc":desc, "message":
message}
```

The string *html_page* is almost a static HTML page but with a few variables. The value for these variables is substituted in the print statement. Point your browser to `http://localhost/cgi-bin/form.py`.

You can enter the data, but not much happens. When you display the form initially, there will be no data in the form. The same program is being called when the form is 'actioned'. So, your code has to be able to differentiate between the two states. You will need to make use of the *cgi* module and add some Python code.

Adding logic

Add the following code after the Python code above:

```
import cgi
def process_form(email, title, desc):
    return "Your entry has been processed"
form = cgi.FieldStorage()
email = form.getvalue("email", "")
title = form.getvalue("title", "")
desc = form.getvalue("desc", "")
if not email or not title:
    message = "Please enter Title and Email"
else:
    message = process_form(email, title, desc)
    email = title = desc = ""
# present the form
print content_header
print html_page % {"title":title, "email":email, "desc":desc, "message":
message}
```


The *cgi* module does not differentiate between *get* and *post* actions. The *getvalue* method picks up the value from the form. The second parameter passed to this method assigns a default value in case the value is not available. The method *process_form* currently does nothing. It could contain the code for storing the data in a database, or whatever else you may need.

Adding some style

However, the impression may be that pages generated through Python are plain and ugly. We can make them even uglier, but colourful! The look and feel will be controlled by a style sheet exactly like the usual HTML pages. So, replace the `<head></head>` line above by:

```
<head>
<style type="text/css">
p {margin-left: 50px}
body {background-color: tan}
input {background-color: yellow}
button {background-color: orange}
</style>
</head>
```

The page now will be more colourful. Incidentally, the styles can usually be read from a CSS file.

WSGI

Now that you know how to write a simple application for the Web, you may want to write an application server in Python. Chances are you would use one like Django, Turbo-Gears, Pylons, etc. But you can use the Web server gateway interface as well, which is supported by various application servers. You can find out more about it at www.wsgi.org or just by searching for 'wsgi tutorials', e.g., http://webpython.codepoint.net/wsgi_tutorial. Try the following simple code in a file *test_app.py*:

```
#!/usr/bin/python
from wsgiref.simple_server import make_server
def application(environ, start_response):
    keys_of_interest = ['PATH_INFO', 'QUERY_STRING', 'REQUEST_METHOD']
    env_vars = (key + ':' + value for key, value in environ.items()
                if key in keys_of_interest)
    msg = '\n'.join(env_vars)
    status = '200 OK'
    response_headers = [('Content-Type', 'text/plain')]
    start_response(status, response_headers)
    return ["Environment Variables of Interest\n\n", msg]
# Start the server and handle just one request
httpd = make_server('localhost', 8080, application)
httpd.handle_request()
```

To create a Web server, you will need to have a method called *application* (not mandatory, but

preferable). The method requires two parameters, *environ* and *start_response*, and returns a list of strings to be displayed. The first parameter is a dictionary containing various environment variables. Your program displays a few of them related to the Web page. The second parameter is a method you call to, well, start the response for the request.

You can convert your code into a Web server using the simple reference server included in Python. Your server would normally keep handling requests and instead of calling *handle_request*, you would call *serve_forever*.

You may notice the use of `()` instead of `[]` to define *env_vars*, using list comprehensions. The use of square brackets creates a list while the use of parentheses creates a generator for the list. Using a generator saves the creation of a temporary list, which can be especially beneficial when dealing with potentially large lists.

You would run the application from the command prompt, *python test_app.py*, and by pointing the browser to *localhost:8080/anything?anyvar=anyvalue*. The result would be as follows:

Environment Variables of Interest

```
REQUEST_METHOD:GET
PATH_INFO:/anything
QUERY_STRING:anyvar=anyval
```

You can change the CGI application you wrote to work with WSGI instead. Your file, *application.py*, should then include the following:

```
#!/usr/bin/python
from wsgiref.simple_server import make_server
from cgi import parse_qs
def application(environ, start_response):
    path_info = environ['PATH_INFO']
    if path_info.endswith('form.py'):
        # for post method, params in a file object
        req_size = int(environ['CONTENT_LENGTH'])
        params = parse_qs(
            environ['wsgi.input'].read(req_size))
        response = process_form(params)
    else:
        response = display_form()
    status = '200 OK'
    response_headers = [('Content-Type', 'text/html')]
    start_response('200 OK', [('Content-Type', 'text/html')])
    return [response]
httpd = make_server('localhost', 8080, application)
httpd.serve_forever()
```

The core application code is not very different from the earlier one. The form responds with a *POST* method and calls *form.py* because the form had called it in the

CGI version; and you too are using the same HTML page and you call *process_form*. For all other URIs, you display a blank form. A URI does not correspond to a file or a method name any more. You can use *path_info* and the parameters returned to perform the actions required for the Web application.


The parameters returned by the *GET* method are in the *QUERY_STRING* environment variable. However, the *POST* method may return large amounts of data; so it is available in a file object referred by *wsgi.input*. The method *parse_qs* parses a string and converts the parameters into a dictionary. Following is a minimal example of *display_form* and *process_form* that you may write:

```
html_page = defined as above
def display_form(title="", desc="", email="", msg=""):
    return html_page % {'title':title, 'desc':desc,
        'email':email, 'message':msg}
# the logic
def process_form(params):
    email = params.get("email",[""])[0]
    title = params.get("title",[""])[0]
    desc = params.get("desc",[""])[0]
    if not email or not title:
        return display_form(title=title, desc=desc, email=email,
            msg="Please enter Title & Email")
    else:
        # the code to save or process the data would appear here
        return display_form(msg="Your message has been processed")
```

The data returned by *POST* is a list of values because a form may include multiple rows with the same field name. So, you pick the first value from the list.

The Web server frameworks make it much easier to handle URLs, manage complex Web pages and associate form fields with database columns. So, you would normally use one of the frameworks. Managing classifieds would be a pretty simple task for any of the frameworks mentioned.

It is hoped that by the wider adoption of WSGI, it will be possible to write components that may be 'callable' from various frameworks or Web applications. Here's an example of plugging Trac within a Pylons framework: wiki.pylonshq.com/display/pylonscookbook/A+Pylons+Controller+with+Trac+as+WSGI+Callable

So, you can easily extend your classifieds application with a wiki, using moin-moin and a fault tracking application using Trac!  **END**

By: Dr Anil Seth

The author is a consultant by profession and can be reached at seth.anil@gmail.com

PLEXUS

Hardware & Networking

Courses Offered

RHCT

RHCE



- Excellent Infrastructure.
- Certified Faculties.
- New batch starts on every Monday.
- Register for RHCE exam and avail **5 DAYS FREE** intensive examination coaching .

Plexus Software Security Systems Pvt.Ltd.

#31/14, Burkit Road, T-Nagar
Chennai-600 017.

Ph:044-2433 7355 Mob: 99401 66510

EMail:training@plexus.co.in

URL: www.plexus.co.in

Other Technical Partners

Authorized Cisco Training



Global Knowledge



A woman in a dark suit stands in the center of a long, brightly lit server room. She is looking at a laptop on a small stand. The room is filled with rows of blue server racks on both sides, and the floor is a light-colored tile. The perspective is looking down the length of the aisle, creating a sense of depth.

Building a Highly-available

Web Server Cluster

Getting started with nginx.

*n*ginx (pronounced as 'engine x') is a powerful HTTP Web server/reverse proxy and IMAP/POP3 reverse proxy.

According to a survey conducted in December 2008 by Netcraft, nginx has grown significantly and has surpassed lighttpd (also known as Lighty). Because of its small memory footprint and high scalability, nginx has found tremendous usage in virtual private servers (VPS).

A reverse proxy is a front end to one or more Web servers. All connections originating from the Internet and destined for Web servers (behind the reverse proxy) are routed through the proxy. Based on the configuration, the proxy may decide to serve the request itself or pass it either partially

or totally to one of the member Web servers behind it. In this manner, the reverse proxy presents a single interface for a set of servers, to the caller. A reverse proxy can thus work as a load balancer distributing the incoming load to member servers. It can also cache the contents being served.

The architecture: A layered approach

I typically use Debian or RHEL in my servers and try to stick with the packages available in the base distribution as far as possible.

For the purpose of this article, I will use four Debian Etch servers and will divide the entire set-up into two layers. The first layer (Layer 1) will have a pair of highly-available (using Heartbeat) *nginx* reverse proxy

installations that will be used to load balance the Web servers located in the second layer (Layer 2). In Layer 2, I will set up two *nginx* installations that serve websites, including PHP pages. Instead of two, *you* can of course have as many Web servers as required. The number of Web servers should be dependent on the total load on the servers. If you feel that the load on the servers is increasing, you can easily add another Web server to the cluster.

In addition to this, there can be a third layer (Layer 3) of database servers. The database is typically used by the applications running on the Web/application servers, where the application directly makes a call to the database using an appropriate method. Database clustering has a slightly different approach and we will take up this subject at a later date. Also, the database layer is independent of our current configuration, and can be added and removed at will without impacting the current set-up.

The installation method of *nginx* is identical on all the four servers in our set-up. From the configuration perspective, the Layer 1 reverse-proxy servers will have identical active-passive configurations and the Layer 2 Web servers will have identical configurations.

The Debian Etch has quite a dated version of *nginx*, but it is good enough for our purposes. So I have decided to stick with this old but stable version.

In order to install *nginx* in Debian Etch, issue the following command as the root user:

```
root@all-servers # apt-get install nginx
```

Configuring *nginx* on Web Servers (Layer 2)

We will start by configuring *nginx* Web servers located in Layer 2 and configuring PHP5 support in it using FastCGI. I will configure *nginx* on one Web server (server1) and the second Web server (server2) will have an identical configuration. We look at the configuration files for server2 later.

Let us first back-up the default configuration file for *nginx*:

```
root@server_1 # cp /etc/nginx/nginx.conf /etc/nginx/nginx.conf.orig
```

Now, create a new */etc/nginx/nginx.conf* file with the following data:

Round-robin DNS

Round-robin DNS is a technique to balance load across multiple servers by providing multiple IP addresses in response to a request for domain resolution. This is typically done by creating multiple A records in the DNS server for a domain. The actual load balancing depends on how the client responds to the response returned by the name server. When a client requests for name resolution, the DNS server responds with a list of redundant IP addresses. Sometimes the resolver tries to arrange the IP address list to give priority to the numerically closer network. A few clients pick the first IP address from the list and a few others may try out alternate addresses in case the first one fails.

RR DNS should not be used as the only method of load balancing as it suffers from a drawback of IP address caching and reuse, both within the DNS hierarchy and also at the client site. While responding to a request the DNS server does not consider the geographically nearer location, network congestion, server load and transaction time, etc.

This is best used when you have uniformly distributed data centres across various geographies. This is also typically used to balance the traffic across multiple data centres within the same geography.

```
user www-data;
worker_processes 1;

error_log /var/log/nginx/error.log;
pid /var/run/nginx.pid;

events {
    worker_connections 1024;
}

http {
    include /etc/nginx/mime.types;
    default_type application/octet-stream;
    access_log /var/log/nginx/access.log;
    sendfile on;
    keepalive_timeout 65;
    tcp_nodelay on;
    gzip on;
```

**Enterprise Mail Server, Linux SBS Server,
Anti SPAM, Antivirus and HTTP Filtering.**

**Bandwidth Management, Internet Access Control,
Content Filtering, Web Access Reporting.**

Customised Linux Product Development.

TechnoInfotech™

1, Vikas Permisses, 11 Bank Street, Fort, Mumbai, India - 400 001. Tel.: 91-22-6633 8900 Ext. 324. info@technoinfotech.com

```
include /etc/nginx/sites-enabled/*;
}
```

We now need to create two directories inside `/etc/nginx/`, namely `sites-available` and `sites-enabled`, as follows:

```
root@server_1 # mkdir /etc/nginx/sites-available /etc/nginx/sites-enabled
```

Let's now create a new file `/etc/nginx/sites-available/default` with the following data:

```
server {
    listen 8001;
    server_name server_1.unixclinic.net;
    access_log /var/log/nginx/server_1.unixclinic.net-access.log;
    error_log /var/log/nginx/server1.unixclinic.net-error.log;

    location / {
        root /var/www;
        index index.html index.htm index.php;
    }
}
```

Subsequent to this, execute the following commands:

```
root@server_1 # cd /etc/nginx/sites-enabled
root@server_1 # ln -s ../sites-available/default
```

Now let us create a simple "Hello World!" HTML file in the document root:

```
root@server_1 # echo "Hello World!" >/var/www/index.html
```

Before getting started, we need to test the *nginx* configuration, and only then start it:

```
root@server_1 # nginx -t
root@server_1 # invoke-rc.d nginx start
```

By accessing the website from the browser by visiting `http://server_1.unixclinic.net` we can verify that the Web server is working fine.

Adding PHP support to *nginx*

Of course, the first step is to install *php5-cgi* packages:

```
root@server_1 # apt-get install php5-common php5-cgi
```

nginx does not have in-built support for *fastcgi* processes unlike Apache and *lighttpd*, and so we have to take charge for managing *fastcgi* processes.

As recommended on the *nginx* wiki, I would like to use the *lighttpd's* *spawn-fcgi* program for *php5-fastcgi* implementation. There are loads of other possibilities that you can check by visiting the URLs mentioned in the resources section at the end of this article.

To get *spawn-fcgi*, I downloaded the latest stable *lighttpd* and compiled it. Once the compilation was done and I had the script, I removed all the build tools.

```
root@server_1 # apt-get install build-essential libpcre3-dev zlib1g-dev
root@server_1 # cd /root
root@server_1 # wget http://www.lighttpd.net/download/lighttpd-1.4.19.tar.gz
root@server_1 # tar -xvzf lighttpd-1.4.19.tar.gz
root@server_1 # cd lighttpd-1.4.19
root@server_1 # ./configure
root@server_1 # make
root@server_1 # cp src/spawn-fcgi /usr/local/bin
```

Now let us remove what we have installed for building the *lighttpd*.

```
root@server_1 # dpkg --purge libpcre3-dev libpcrecpp0 \
build-essential cpp cpp-4.1 g++ g++-4.1 gcc gcc-4.1 \
libc6-dev libssp0 libstdc++6-4.1-dev linux-kernel-headers \
zlib1g-dev libbz2-dev root@server_1 # rm -rf /root/lighttpd-1.4.19
```

We can start the *fast-cgi* server as follows:

```
root@server_1 # /usr/bin/spawn-fcgi -a 127.0.0.1 \
-p 9000 -u www-data -f /usr/bin/php5-cgi
```

This can be verified by:

```
root@server_1 # ps aux | grep php
```

To start the *fastcgi* server after every reboot, we can put the following line in `/etc/rc.local` file, or as suggested on the *nginx* wiki, we can write a custom script also. I have chosen to take the ancient route of `/etc/rc.local` file:

```
/usr/bin/spawn-fcgi -a 127.0.0.1 -p 9000 -u www-data -f /usr/bin/php5-cgi
exit 0
```

Now, to configure *nginx* to pass all incoming requests for PHP files to the *fastcgi* process listening on port 9000, we need to add the following location directive to `/etc/nginx/sites-available/default` file. This line needs to be added before the closing of the server directive:

```
location ~ \.php$ {
    include /etc/nginx/fastcgi_params;
    fastcgi_pass 127.0.0.1:9000;
    fastcgi_index index.php;
    fastcgi_param SCRIPT_FILENAME /var/www/$fastcgi_script_name;
}
```

Finally, let us create a simple PHP script in the document root to test this. The traditional route is to create a *phpinfo()* script, but I will just create a "Hello World!" script for security reasons. Open the `/var/www/`

index.php file in a text editor and enter the following line:

```
<? echo "Hello World! This is server 1." ?>
```

Visiting the website http://server_1.unixclinic.net/index.php should verify the working of PHP.

Configuration of Server2

Before we set up load balancing using the reverse proxy feature of *nginx*, we need to have a second identical server. Although as an experiment, you can configure another virtual server in *nginx* listening on a different port and use it, we will use another server that is configured identically to this primary server. The following is the virtual server configuration of server2, which is a separate physical server:

```
root@server_2 # cat /etc/nginx/sites-enabled/default
server {
    listen 8001;
    server_name server_2.unixclinic.net;
    access_log /var/log/nginx/server_2.unixclinic.net-access.log;
    error_log /var/log/nginx/server_2.unixclinic.net-error.log;

    location / {
        root /var/www;
        index index.html index.htm index.php;
    }

    location ~ \.php$ {
        include /etc/nginx/fastcgi_params;
        fastcgi_pass 127.0.0.1:9000;
        fastcgi_index index.php;
        fastcgi_param SCRIPT_FILENAME /var/www$fastcgi_script_name;
    }
}
```

The content of the *index.php* file in the document root is as follows:

```
root@server_2 # cat /var/www/index.php
<? echo "Hello World! This is server 2." ?>
```

Configuring the reverse proxy (Layer 1)

After getting the member servers ready, we will now proceed to configure *nginx* as the reverse proxy. Open the */etc/sites-available/rev-proxy-lb* file in a text editor and enter the following data:

```
upstream web_servers {
    server server_1.unixclinic.net:8001 max_fails=2 fail_timeout=30s;
    server server_2.unixclinic.net:8001 max_fails=2 fail_timeout=30s;
}

server {
    listen 80;
    server_name www.unixclinic.net;
```

Connection distribution methods

The connection distribution methods in *nginx* describe how the connection load is balanced or distributed across member servers. We have already looked at the *weight* parameter in detail, which distributes the connection based on the weight assigned to a member server.

Another method for connection distribution is *ip_hash*. Sometimes it is required that a client request is always transferred to a particular member server—the *ip_hash* directive can be used in such cases. When a client connects to the server, the *ip_hash* directive calculates the hash of the client's IP address and keeps a track of the member server it is connecting to. All subsequent requests from that client are passed to the same member server, if it is available.

In the case of a member server being unavailable, the requests are passed on to another member server. If a certain member server has to be taken down for some reason and for quite some time, then it is advised to mark that server as 'down'. The *weight* and *ip_hash* directives are incompatible with each other and hence cannot be combined together.

```
rewrite ^/(.*) http://unixclinic.net/$1 permanent;
}

server {
    listen 80;
    server_name unixclinic.net;
    access_log /var/log/nginx/rproxy_1-access.log;
    error_log /var/log/nginx/rproxy_1-error.log;

    location / {
        proxy_pass http://web_servers;
    }
}
```

The first server directive just contains a rewrite to redirect all requests coming to <http://www.unixclinic.net> to <http://unixclinic.net>.

The *upstream* directive is in the *nginx_http_upstream* module that balances load across multiple back-end servers. This module uses a simple round-robin load-balancing algorithm. The *upstream* directive specifies a set of servers that can be used in other directives such as *proxy_pass* and *fastcgi_pass*.

The server directive specifies the name of the member server and the parameters applicable for a server.

- The name part can contain a domain name, IP address, a port number and a UNIX socket. If a domain name resolves to multiple IP addresses (multiple A records in the DNS for a domain, see note), then all the IP addresses are used.
- A weight can be assigned to each server to specify its priority level for handling requests. If a weight is not assigned, then it is considered as 1. For example, if a weight of 2 is specified for *server_1.unixclinic.net* as follows, then for every three requests, two will be

passed to server_1 and one to server_2.

```
upstream web_servers {
    server server_1.unixclinic.net:8001 weight=2;
    server server_2.unixclinic.net:8001;
}
```

- The *max_fails* parameter specifies the maximum number of failed connection attempts with a member server within a specified time period. This time period is specified by another parameter *fail_timeout*. The default value of this parameter is 1, and if this is set to 0 then the check is disabled. Setting this to 0 is certainly not recommended and I would advise setting this to at least 2 or 3.
- The *fail_timeout* parameter contains the time duration in seconds. If all connection attempts with a member server for the specified time fail, then this would count as one failure. In our configuration above, the *fail_timeout* value is 30 seconds. This means that if the reverse proxy's connection attempts with one or more of our member servers failed for 30 seconds, then it would count as failed once, and another 30 second failure consecutively will take the 'failed' count to two. After the failed count reaches two, which is what is the value of *max_fails* parameter in our case, the member server is marked as unresponsive for a certain amount of time. The default value for this is 10 seconds. And this timeout is controlled by other directives called *proxy_connect_timeout* and *proxy_read_timeout*. Read the *nginx* wiki for more details about this.
- The *down* parameter marks the member server as permanently down. This is typically used with the directive *ip_hash* (discussed below).
- The *backup* parameter is only available in *nginx* version 0.6.7 or later. So if you decide to compile the latest *nginx* version for your use then this could be used. This parameter specifies the member server as a back-up server in case all the other member servers are busy or down.

The Debian Etch has *nginx* version 0.4.13-2, whereas the 'backports' repository has *nginx* version 0.5.35-1. If you have decided to configure the 'backports' repository in your server and then install *nginx*, you can make use of a new feature of the upstream directive, which is available since version 0.5.18. This feature allows you to log a few additional variables via the log module. These variables are:

- *\$upstream_addr*—address of the upstream server that handled the request
- *\$upstream_status*—upstream server status of the answer
- *\$upstream_response_time*—this is recorded in milliseconds. Several answers will be divided by commas and colons.
- *\$upstream_http_\$HEADER*


Finally, to activate the reverse proxy, issue the following commands:

```
# cd /etc/sites-enabled
# ln -s ../sites-available/rev-proxy-lb
# nginx -t
# invoke-rc.d nginx reload
```

After this, visit <http://unixclinic.net/index.php>. Each time you access it, you should notice that the pages are being served according to the weight and from different servers. I have specified *index.php* explicitly because *index.html* has a higher priority on this, and if you remember, we have created *index.html* above, while testing the Web server.

Bottlenecks in this set-up

The websites are redundant and load balanced according to the simple algorithm that *nginx* provides. There are a few bottlenecks in this set-up, which can cause a hurdle in providing a true highly available set-up. The first bottleneck that we can easily see is that we have only one reverse proxy server and in case this goes down, our website too will be down. In order to resolve this we will need to set up another instance of a reverse proxy server as a secondary server, which will take control of the domain(s) being served in case the primary load balancer goes down. In other words, we will need to set up an active-passive clustering between the two *nginx* reverse proxy servers.

We will take up this and some other issues in a subsequent article to see how we can achieve a better redundancy.  **END**

Resources

Nginx and Php5 Fast-cgi

- *nginx* FastCGI example: wiki.codemongers.com/NginxFcgiExample
- *nginx* tips: dominiek.com/articles/2007/5/18/bye-redskin-zdravstvujte-redstar-nginx-tips
- *nginx*, PHP and a PHP FastCGI daemon init script: blog.codefront.net/2007/06/11/nginx-php-and-a-php-fastcgi-daemon-init-script
- How to set up Ubuntu with *nginx*/php/mysql/ruby/radiant/mongrel/monit and a simple firewall: johnmuhl.com/notebook/server
- *Nginx* HTTP Server/PHP5 (with FastCGI and xcache) on Ubuntu: www.howtoforge.com/nginx_php5_fast_cgi_xcache_ubuntu7.04

Load Balancing

- Round-robin DNS: en.wikipedia.org/wiki/Round_robin_DNS
- How to configure round-robin and load-balancing: www.zytrax.com/books/dns/ch9/rr.html
- Load balancing Web applications: www.onjava.com/pub/a/onjava/2001/09/26/load.html

By: Ajitabh Pandey

The author has more than 12 years of diversified IT industry experience in training, support and consulting. You can know more about him at ajitabhpandey.info and shoot him an e-mail at ajitabhpandey@ajitabhpandey.info

LINUX JOBS

Post: Testing Lead/ QA Lead

Company: Agilis Information Technologies Intl. Pvt. Ltd.
Profile: On hands experience in testing of products or modules in Test Lead Role is desired alongwith an experience in querying the Oracle database.
Exp: 4-7
Location: Delhi/NCR
Email: careers@agilisinternational.com

Post: Project Lead- Java

Company: Telecom Product Multinational
Profile: Would be responsible for the delivery of the project module from technical perspective, analysis, design and estimation of work.
Exp: 5-7
Location: Bengaluru
Email: rinku@bestitquest.com

Post: Oracle DBA (RAC)

Company: Ismart Panache Inc
Profile: Hands on experience of Administration in Oracle (Production) environment 9i/10g/RAC with good knowledge of UNIX/ Solaris/ Linux (RHEL)/ Windows.
Exp: 4-8
Location: Gurgaon
Email: adeshpande@ismartpanache.com

Post: Sun/ Storage Engineer

Company: Gulf Computers L.L.C
Profile: Responsibilities include installation of Solaris on sun servers, updating recommended patches and all the other patches on Solaris, configuring system and network services as per the user requests.
Exp: 4-7
Location: Dubai/ UAE
Email: ranjini@gulfcomputers.ae

Post: Linux System Admin

Company: New Horizons
Profile: Required candidate should be B.E/ B.Tech, Diploma with 5-8 years experience.
Exp: 5-8
Location: Noida
Email: Test_naukri@yahoo.com

Post: Solaris Admin

Company: Tech People Consulting, Inc
Profile: Minimum of 8-10 years of experience as a UNIX systems admin with very strong Linux systems admin experience is desired.
Exp: 6-8, Location: United States (U.S)
Email: usjobs@tpc-jobs.com

Post: Oracle Apps DBA

Company: Bahwan Cyber Tek
Profile: Should have 3+ yrs of relevant experience, must have good exposure on OS as well as instance maintenance, should have worked for atleast 2 implementation projects.
Exp: 4-9, Location: Chennai
Email: careers@bahwancybertek.com

Post: ITIL Professionals

Company: Infinite Computer Solutions India Pvt Ltd
Profile: Should be certified and worked on ITIL for more than 1 year.Knowledge of ITIL skills are mandatory.
Exp: 2.5 -4
Location: Bangalore
Email: deepaa@infics.com

Post: Sr. Test Lead

Company: Minecode Solutions Pvt. Ltd.
Profile: The Test Manager will be directly involved in ensuring the quality of applications, products and tools for the company. Degree in Computer Science or 8-10+ years related work experience.
Exp: 5-10
Location: Delhi/NCR
Email: hrindia@minecode.com

Post: Technical Support Engineer- Telecom Core Network

Company: Mavenir Systems
Profile: Prior Deployment/ Maintenance/ testing experience in GSM/ IMS/ CDMA/ UMTS Core Network area- MSC/ VLR, HLR/ HSS, SMSC, CSCF etc. is desired. Strong Knowledge on UNIX & Linux is desired.
Exp: 4-7
Location: Bengaluru
Email: gudla@mavenir.com

Post: Lead Product Developer

Company: BMC Software
Profile: 7-10 years experience in Software Quality Engineering with focus on testing enterprise- class, web applications.
Exp: 5-7
Location: Pune
Email: diwakar_rautela@bmc.com

Post: Tech Lead- Internet Technologies

Company: Directi
Profile: 8+ years of software engineering experience, of which 3+ years have been spent leading a team with an experience in developing products in one or more of Java, C#, PHP, Ajax / Javascript.
Exp: 8-12, Location: Mumbai
Email: php-tl-jobs-mum@directi.com

Post: H1B Visa for Java/J2ee Professionals

Company: TechnoGiant Solutions Inc.
Profile: Should have knowledge of Java / J2EE, Unix / Linux, Websphere / WSAD / Servlets, MQ Series / XML with relevant experience in the same.
Exp: 4-9
Location: United States (U.S)
Email: Vinay.Sharma@technogiant.com

Post: EMC Storage Management Systems Administrator

Company: Source One Management Services Pvt.Ltd.
Profile: Ideal candidate should possess a Bachelor's Degree & good Communication skills with Diploma in Computer Science in Electronics. 3+ years must experience.
Exp: 4-9
Location: Chennai
Email: amit.mishra@sourceone.co.in

Post: Sr. SCM Professional

Company: Yahoo Software Development India Pvt. Ltd.
Profile: 7-11 years experience is required.
Exp: 7-11,
Location: Bengaluru
Email: gopald@yahoo-inc.com

Post: Storage Admin

Company: Omnitech Infosolutions Ltd.
Profile: Job requires a BS/ MS in Computer Science, Information Systems, Engineering with at least 6 years track record in consulting/ IT management.
Exp: 7-12
Location: Mumbai
Email: alan@omnitechindia.com

Post: Sr. Specialist Systems Network Administrator

Company: Office Beacon Administrative Services Pvt. Ltd.
Profile: Should have an experience in LINUX, UNIX installation and configuration is a must with knowledge of Apache, Mysql, & PHP.
Exp: 8-13
Location: Vadodara
Email: myadav@officebeacon.com

Post: Web Developer

Company: ITECS Inc
Profile: Must have skills include at least 2 years Adobe Flex 2/3 with ActionScript 3 experience with additional Adobe products and tools including Flex SDK, and ColdFusion at least 5 years of overall experience.
Exp: 5-7
Location: United States (U.S)
Email: careers@itecsus.com

Post: MysqlDBa Linux Admin

Company: Tech People Consulting, Inc
Profile: Atleast 3 yrs of experience in MySQL DBA is desired with knowledge of Linux Administration.
Exp: 5-6
Location: United States (U.S)
Email: usjobs@tpc-jobs.com

Post: HP Unix Admin

Company: IMSI India Pvt Ltd
Profile: Candidate will be desired to monitor concurrent Logins/ Connections, Processor Utilization, Network Utilization related to server with knowledge in SDS Administration.
Exp: 4-6
Location: Mumbai
Email: deepa.aggarwal@imsiglobal.com

Post: Tech Lead - Java/J2ee (Networking)

Company: Photon Infotech Pvt. Ltd.
Profile: Should have knowledge of J2EE, .Net, Linux, UNIX, Windows 2003 server & should be familiar with Routers, Switches, Firewalls configurations, Java, Shell Scripting, PL/SQL, XML, TCL.
Exp: 6-11
Location: Bengaluru
Email: rajagopal.s@photoninfotech.com



Get the Best Paying Jobs @ naukri.com
2,00,000 Jobs to choose from!

35,000 Recruiters | Get Jobs by email | Block Your Current Employers

Powered by

naukri.com
India's No.1 Job Site

For complete recruitment solutions Employers can call us at 0120-(3082097) or email us at vinny.ganju@naukri.com



What's in the Glass(Fish)?

Part 2: Getting Started With the App Server

Last month, we discussed application servers, and in particular, the GlassFish app server. This month we will look at how to get started with GlassFish version 2 on Debian GNU/Linux. For other distributions, the procedure will be almost the same.

Installation of GlassFish

GlassFish v2 requires JDK 1.5 or a later version. So if you are using Fedora 9 or Ubuntu 8.4, you already have OpenJDK 6. You can check the version of Java with the `java -version` command at a terminal prompt. This command will print the available version of Java on your system. For example, on my system the available version of Java is 1.4.2 as can be seen from the following terminal output:

```
mca05@debian:~$ java -version
java version "1.4.2"
gij (GNU libgcj) version 4.1.2 20061115 (prerelease) (Debian 4.1.1-20)
```

Copyright (C) 2006 Free Software Foundation, Inc.

This is free software; see the source for copyright conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

So I have to install JDK 1.5 or later before installing and configuring GlassFish. Verify that your `/etc/apt/sources.list` file has the following entries and you are connected to the Internet:

```
deb http://security.debian.org/ etch/updates main contrib non-free
deb http://ftp.debian.org/debian etch main contrib non-free
deb-src http://ftp.debian.org/debian etch main contrib non-free
```

Then simply run the following command:

```
apt-get install sun-java5-jdk
```

It will install JDK 1.5 or above on your system. Now that we

have successfully completed the basic requirements, download the GlassFish jar file from download.java.net/javaee5/v2ur2/promoted/Linux/glassfish-installer-v2ur2-b04-linux.jar and issue the following command:

```
java -jar -Xmx=256 glassfish-installer-v2ur2-b04-linux.jar
```

This command will show a licence agreement window—it's dual licensed under CDDL and GPL. After accepting the licence it will create a directory named *glassfish* in your current working directory. After the installation is complete, we need to set it up before being able to use it.

Configuration and set-up

To set up GlassFish navigate to the *glassfish* directory from the terminal. Here you will see a file named *setup.xml*. This file is the input file for ant builder scripts that will configure GlassFish. *setup.xml* contains the following configuration property elements:

```
<property name="install.home" value="${basedir}"/>
<property name="domain.name" value="domain1"/>
<property name="instance.name" value="server"/>
<property name="admin.user" value="admin"/>
<property name="admin.password" value="adminadmin"/>
<property name="admin.port" value="4848"/>
<property name="instance.port" value="8080"/>
<property name="orb.port" value="3700"/>
<property name="imq.port" value="7676"/>
<property name="https.port" value="8181"/>
<property name="glassfish.license" value="LICENSE.txt"/>
<property name="glassfish.zipfile" value="../glassfish-image-pe.zip"/>
<property name="glassfish.release.name" value="glassfish"/>
<property name="glassfish_pe.release.name" value="glassfish-pe"/>
<property name="glassfish.class" value="glassfish.class"/>
<property name="persistence-location" value="${install.home}/../glassfish-persistence"/>
<property name="glassfish-persistence.zipfile" value="glassfish-persistence-image.zip"/>
<property name="glassfish_persistence.class" value="glassfish_persistence.class"/>

<property name="jarpack-task.jar" value="${install.home}/lib/Pack200Task.jar"/>

<property name="adminpassfile" value="${install.home}/passfile"/>
<property name="asadmin.prefs.file" value="${install.home}/config/asadminenv.conf"/>
<property name="asadmin.default.profile" value="developer"/>
<property name="asadmin.default.secure" value="false"/>
```

Before setting up the server, you need to make some changes in the property values according to your preferences. First of all, change the default *admin.password* property, which by default is set as 'adminadmin' (see line 5 in the above snippet). You

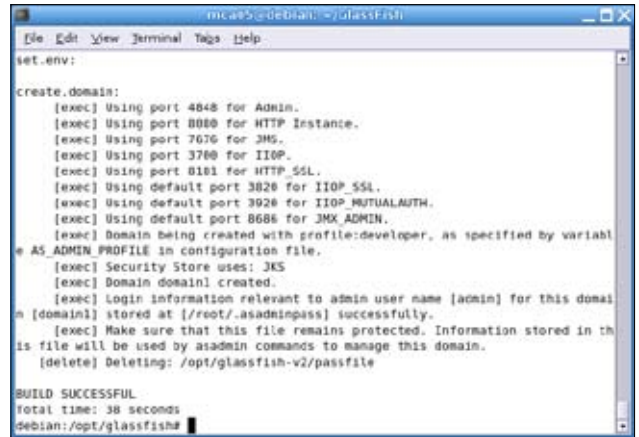


Figure 1: GlassFish build successful

need to remember this password, as it will be used for login and deploying applications later. You can also change the ports on which the admin and applications will listen. Change them accordingly if they conflict with services that listen on the same ports.

After making changes in the configuration file, it's time to start the build process. I will recommend a set-up process such that in the future you can easily update GlassFish with its new releases.

We will move *glassfish* to */opt/glassfish-v2*, and then create a soft-link (*ln -s /opt/glassfish-v2 /opt/glassfish*) such that in the future we can easily upgrade it. Make the ant builder script executable by issuing the following command:

```
chmod u+x /opt/glassfish/lib/ant/bin/ant
```

Now, to start the build process, issue the following command:

```
/opt/glassfish/lib/ant/bin/ant -f setup.xml
```

You should soon get a 'build successful' message as shown in Figure 1.

To start the server and check whether it is installed and configured successfully, issue the following command:

```
/opt/glassfish/bin/asadmin start-domain
```

...and point your browser to *http://localhost:4848*. You should get a login page as shown in Figure 2. If you don't, then recheck the whole procedure for the possible error.

For login, use the password you have applied for the *admin*. password property earlier.

Developing a simple Web application

In this section I will develop a simple dynamic Web application based on the Model View Controller (MVC) design pattern, which will generate prime numbers for its clients. An MVC design pattern divides an application into the following components:



Figure 2: App server login page

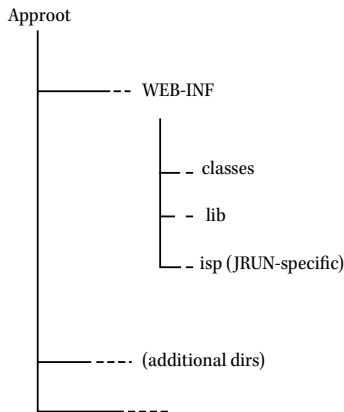


Figure 3: Web application directory structure

1. Model: The content or the actual data (or logic)
2. View: The presentation of data to the clients
3. Controller: It controls the application—for example, calling the model for business logic and then forwarding the request to 'view' for its presentation.

The Web application will be based on servlets and JSP. In very simple terms, servlets are server side applications that can dynamically extend the functionality of a Web server, and JSP is a presentation technology based on servlets. If you want to know more details about servlets and JSP then I'd recommend doing a Google search for the answers.

Typical Web application development on an app server requires a directory structure like in Figure 3. Here, *Approot* is the root directory for our Web application. It contains the default page that will be served by the server to its clients.

In addition, it contains a special directory named *WEB-INF*. This is the directory where our Web application-specific files will be stored. *WEB-INF* contains the following files and folders:

- The deployment descriptor file named *web.xml*.
- A *classes* folder that contains the model and controller classes of our MVC Web application
- A *lib* folder that contains third-party jar files
- A *jsp* folder that is JRUN specific; it will not be used in our sample application.

For our prime number generator Web application, create a directory structure as shown below:

- The top-level directory *Approot* contains the default page *index.html*, the *img* directory (which contains the images used in *index.html*), *result.jsp* (view of our MVC Web application), and the *WEB-INF* directory.
- *WEB-INF* contains *classes* and the *web.xml* file (a

deployment descriptor file used by the server to map servlet names to its classes).

There is no need of *lib* and *jsp* directories, as we are not using third-party *jar* files or *jrnl*.

The default page for the prime number generating Web application, *index.html*, looks like what's shown below:

```

<html>
<!-- Some sample html-->
<head>

        <title>Prime number generator Web Service</Title>

</head>
<body>

        <h1 align="center">Prime number generator Web Service</h1>

<center>

</center>
<center>

        <form method="GET" action="PrimeSearcher.do"
enctype="multipart/form-data">

        <b>Enter the number till which you need prime numbers:</b>
        <input type="text" name="number" size="5"/><br/>
        </center>
        <br>
        <center>

<input type="SUBMIT" value="GENERATE PRIMES"/>
</center>
<br>
<br>
<br>
</form>
</body>
</html>
  
```

This *index.html* is presented to clients who access the Web application. As you can see in the above snippet, the method is *GET* and the action servlet is *PrimeSearcher.do*. *PrimeSearcher.do* is just a name given to our controller servlet and *.do* is just a convention. The *GET* method indicates that whenever a client clicks the *GENERATE PRIME* button, the *doGet* method of the servlet is called.

The deployment descriptor file, *web.xml* looks like what's given below:

```

<web-app xmlns="http://java.sun.com/xml/ns/j2ee"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://java.sun.com/xml/ns/j2ee/web-app_2_4.xsd"
version="2.4">
<servlet>

        <servlet-name>PrimeSearcher</servlet-name>
        <servlet-class>Controller</servlet-class>
</servlet>

<servlet-mapping>
        <servlet-name>PrimeSearcher</servlet-name>
  
```



```

        <url-pattern>/PrimeSearcher.do</url-pattern>
    </servlet-mapping>
</web-app>

```

The deployment descriptor file *web.xml* is specific to a Web application, and the app server uses it to map servlet names to its classes. Each Web application has its own *web.xml* file. It maps the servlet names to servlet classes—for example, in our application, *PrimeSearcher.do* is mapped to *PrimeSearcher*, and *PrimeSearcher* is mapped to *Controller* class.

The ‘view’ part of our Web application is a JSP page (*result.jsp*) that generates output for its client. It looks like the following:

```

<%@ page import="java.util.*"%>
<html>
<body>
<h1 align="center">Prime Number generator View</h1>
<p>
<%
    Iterator it= (Iterator)request.getAttribute("primes");
    out.append("primes are:"+"<br>");
    while(it.hasNext())
    {
        out.println(""+it.next()+"<br>");
    }
%>
</body>
</html>

```

This *result.jsp* is internally translated into a servlet by the container. The JSP gets the attribute named *primes* attached to the *HttpRequest* object by the *Controller*. This *primes* attribute is an iterator that contains the prime numbers generated by the model class.

The *Controller* code is shown below:

```

/*****controller.java*****/
import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

public class Controller extends HttpServlet {

    public void doGet(HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException {

        Long num= new Long(req.getParameter("number"));
        Iterator it= model.generatePrime(num.longValue()).iterator();
        req.setAttribute("primes", it);
        RequestDispatcher view= req.getRequestDispatcher("view.jsp");
        view.forward(req, res);
    }

    private Model model = new Model();
}

```

```

}

```

It is a simple servlet class that extends *HttpServlet*. It takes the help of the *Model* class to generate the prime numbers. This *Model* class contains the actual logic to generate primes. The code looks like this:

```

/*****Model.java*****/
import java.util.ArrayList;
import java.util.Iterator;
import java.util.List;

public class Model {

    List generatePrime(long l)
    {
        long number=2;
        List primes= new ArrayList();
        primes.add(number);
        long sqrt ;
        while(number<l)
        {
            number++;
            if(checkPrime(number)==true)
            {
                primes.add(number);
            }
        }
        return primes;
    }

    boolean checkPrime(long l)
    {
        long sqrt=(long) Math.sqrt(l);
        for (long i = 2; i <= sqrt; i += 1) {
            if (l % i == 0)
                return false;
            else
                ;
        }
        return true;
    }
}

```

This model class contains the actual logic to generate primes, and returns a list of generated prime numbers to the *Controller*.

To understand the entire functionality, assume that when a client requests our Web application to generate primes, then the request will be transferred to the controller class that gets the number entered by the client from the *HttpRequest* object. This number is passed to a model class, which actually generates prime numbers and has the real logic of our application. The model class returns a list that contains prime numbers. The *Controller* class gets the iterator of this list and attaches this iterator as an attribute to the *HttpRequest* object, and forwards the remaining functionality to the ‘view’ of our




Figure 4: The applications branch on the app server

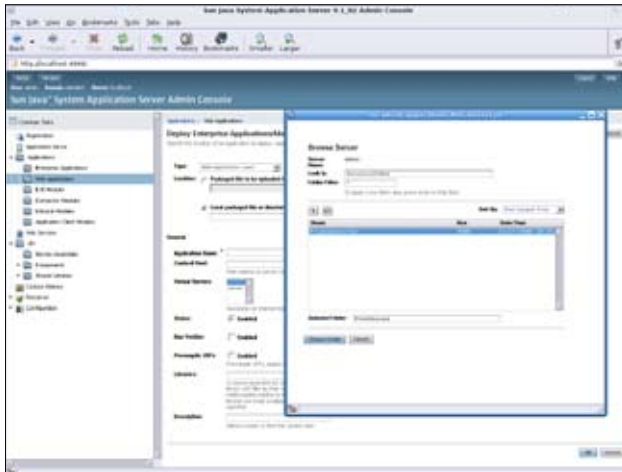


Figure 5: Browsing the locally packaged file to deploy on the Web server




Figure 6: A local application deployed on the app server

application, which is a JSP file. 'View' gets this iterator object and generates a page containing the generated prime numbers.

Deploying the Web application on GlassFish

First of all, compile the Controller and Model classes using *javaee.jar* in *classpath*. The GlassFish *lib* folder contains this *javaee.jar*, which in turn contains the *javax.servlet* api.

After compiling the controller and model Java files, it's time to deploy the Web application. First copy the controller and model *.class* files in the classes folder of the *WEB-INF* directory. Then start the GlassFish server and log in. Go to the applications branch of the tree as shown in Figure 4. Then go to the Web application branch. Check the radio button for "Local packaged file or directory..." and click the 'browse' button (Figure 5). It will show a 'directory chooser'. Choose the prime generator (folder) that contains our Web application. Then click on the OK button. That's it! You have successfully deployed the application and will see a page like what's shown in Figure 6.

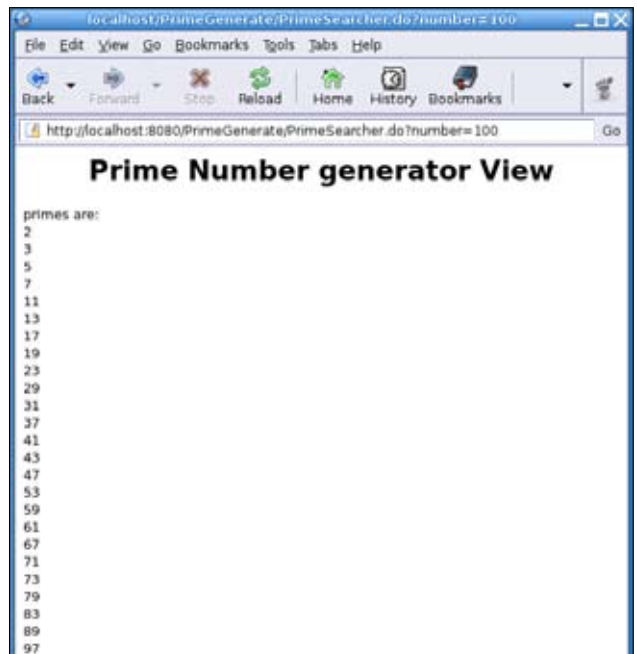


Figure 7: The prime number generator application

Running the Web application

Finally, it's time to run the application. Simply click the launch link (see Figure 6). You will see the default page of our application as shown in Figure 7. Enter the number up to which you need prime numbers to be sorted out. Click the "Generate Primes" button and you will get the primes as shown in Figure 8.

That's it for now and happy development with GlassFish and Java EE! 

By: **Rajeev Kumar**

The author is a software engineer working at Aricent Technologies. He loves working with FOSS and GNU/Linux, as well as Java Enterprise Edition. You can reach him at rajeevcoder@yahoo.com

Which one is your career **goal?**



medical



engineering



computer software



computer hardware



hotel management



clinical research

We help you **decide**

Hundreds of Educational Institutions Offering Professional Courses & Training Under A Single Roof

EduTech's Highlights :


- * An event totally focused on professional education
- * A one-stop facility for students seeking admission into professional courses
- * Counselling by industry top-shots (HR and recruitments)
- * Major plan for advertising and promotion of EduTech Expo 2009
- * Powered by EFY Group's vision to boost professional education
- * Courses catering school pass-outs & fresh graduates

**EDUTECH
EXPO
2009**

13th - 14th JUNE 2009
PRAGATI MAIDAN, NEW DELHI



EFY Enterprises Pvt Ltd,
D-87/1, Okhla Indl. Area, Phase-I, New Delhi-110020;
Ph:011-26810601-03; Fax: 011-26817563



Tips & Tricks



How to use *diff* and *patch*

You have created a program *prog.c* and also copied to *prog.c.old*. You post *prog.c* to users. Next, you make changes to *prog.c* and want to release a patch. Here is how you can do it:

```
diff -c prog.c.old prog.c > prog.patch
```

Users can now get the latest updates by running:

```
patch < prog.patch
```

—Remin Raphael, remin13@gmail.com



Five useful shell-prompt tips

1) We always define the alias for a shortcut. But, sometimes, we don't need that and want the original effect of the command. For example:

```
alias ls='ls -la'
```

Now to restrict the alias effect and see the native command, we can execute it in any of the following three ways:

```
$ command ls
$ \ls
$ "ls"
```

2) If you want to check whether an alias is assigned to a particular command, you can use the *type* command with an alias name as follows:

```
$ type ls
```

You can also use the *unalias* command to revert back to the original *ls* command:

```
$ unalias ls
```

This will display the alias name, if the alias is assigned.

3) To open an application from the command line in GNOME, we can use the *gnome-open* command. For example:

```
$ gnome-open jash.xls
```

This will open the file with the application associated with it.

4) Here is a useful way to use the append operator. To append text at the end of a file, we simply use the *>>* character. But if we want to append text at the beginning of the file, run the following command:

```
$ echo "hi jash" | cat - file.txt > /tmp/out && mv /tmp/out file.txt
```

By using this command, the "hi jash" line will be added at the beginning of the file.

5) If you want to find all files with the extension *.txt* and *.jpg*, then use the following command:

```
$ find . -type f -iname "*.txt" -or -iname "*.jpg"
```

...where the *-type f* parameter is used to find files and *-or* is used for the 'or' operation.

If you want to exclude hidden files in this 'find', then use the *!* (not) operator:

```
$ find . -type f -iname "*.txt" ! -iname ".*"
```

—Jasvendar Singh M. Chokdayat,
theindianjash@yahoo.com



Identify zombies

Do you need to identify the process ID of a zombie process, and also know which program the process belongs to? Execute the following command to get the list of zombie processes:


```
ps aux | awk '{print $8 " " $2}' | grep -w Z
```

You will get an output like what's shown below:

```
Z 6780
Z 8976
```

Here, 6780 and 8976 are zombie process IDs.

You can check the program associated with the process using the *ps* command and kill the process if required.

If you want to display the process that is consuming more resources, here is the command:

```
top | head -n 8 | tail -n 1
```

—Bharathiraja, pbrthemaster@gmail.com



Moving quickly from anywhere, with \$CDPATH

Do you know the environment variable *\$CDPATH*? It lets you define some path to look for directories when moving with the command *CD*. You can define multiple paths in this variable. This can be useful if you have some directories that you access more frequently than others. As this variable defines the order that you move from one directory to another, I suggest you to keep the *~* directory in the first position of your variable declaration. Of course, you can set this variable in your *~/.bashrc*.

```
smart@grimm:~$ export CDPATH=:~/~/
smart@grimm:~$ cd usr
/usr
smart@grimm:/usr$ cd Desktop
/home/nicolas/Desktop
```

—Remin Raphael, remin13@gmail.com



Hardware info on live systems

The *dmidecode --type memory* command is used to know the type of RAM and its speed without opening your computer. It can also be used to get to know details about the BIOS, motherboard, processor, cache, slots, etc. If we want to get information about the BIOS, use the following command:

```
dmidecode --type bios
```

For cache, use:

```
dmidecode --type cache
```

...and so on.

—Elvise James, elvisejames@gmail.com



Create a directory with different permissions

When we need a directory with specific permissions, we generally create the directory using the *mkdir* command and then set the required permissions using *chmod*.

Did you know you define the access permissions while creating the directory itself? Try the following command:

```
[root@server ~]# mkdir /root/test1 -v -m 1777
mkdir: created directory '/root/test1'
```

...where we have used the *-m* flag to define the permission we need to set (in our case it's 1777, that is read, write and execute permissions for the directory are given to everyone), and *-v* stands for verbose mode.

Now, you can check the directory permission by a simple *ls* command as follows:

```
[root@server ~]# ls -lhd /root/test1
drwxrwxrwt 2 root root 4.0K Jul 29 01:42 /root/test1
```

—Govindarajalu govind.rajalu@gmail.com



Save complex documents from any source

Latest versions of GNOME and KDE (along with CUPS) support creating PDF files out of any printable document. This facility can be used to save any printable document as a PDF file.

1. Select the printing option, usually *File→Print...*, in the application hosting the document.
2. Select the PDF option as the printer name in the dialogue box—*Print to File/PDF*.
3. If applicable, type a file name, set output options, and click *Print*.

This process is very useful for quickly saving dynamically-generated Web pages such as e-mails, confirmation forms, reports, etc.

—Saurav Sengupta, saurav.mailbox@gmail.com



Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at www.linuxforu.com. The sender of each published tip will get an LFY T-shirt.



Sandya Mannarswamy

Welcome to another installment of CodeSport. This month's column focuses on computational complexity and the lower bounds for algorithms. In particular, we'll show that any algorithm to find the maximum in an array of N elements has a lower bound of $O(N)$ by using an adversary argument.

Thanks to all the readers who sent in their comments on the problems we discussed in last month's column, in which the takeaway question was a variant of the 3-sum problem. You were given three sets of numbers A , B and C , each containing N numbers. Readers were asked to come up with an algorithm to determine whether there was a triple $a \in A$, $b \in B$ and $c \in C$ such that $a + b = c$? It is quite easy to come up with an $O(N^2 \log N)$ algorithm. But you need to come up with an $O(N^2)$ algorithm for this.

Instead of a separate algorithm for our takeaway problem, we will first look at how we can reduce the 3-sum problem to this. Given a set S of n integers, do there exist three elements $\{a, b, c\} \in S$ such that $a + b + c = 0$? Let us denote the takeaway problem as PR2 and the 3-sum problem as PR1. Since we know that the 3-sum problem has a lower bound of $O(N^2)$, PR2 also is equally hard. Hence, the best we can hope for to solve PR2 is an $O(N^2)$ algorithm. Recall that the reduction from a 3-sum problem to PR2 can be given by $A=S$, $B=S$ and $C=-S$. Hence, $a \in A$, $b \in B$, $c \in C$, and a solution of PR2 which satisfies $a + b = c$, will also satisfy PR1. Now, we would like to look at how we can reduce PR1 to PR2; i.e., given a problem instance of PR1, we want to formulate a problem instance of PR2 that can be reduced to PR1. In that case, we can use the solution of the 3-sum problem to solve PR2.

We create one set S such that whenever three elements in S add up to zero, there are three elements $a \in A$, $b \in B$ and $c \in C$, such that $a+b = c$. For the sake of simplicity, let us assume that all elements are positive. We

define a value m , which is equal to twice the maximum element among all members of A , B and C . We now try and construct set S from the elements of A , B and C as follows: For each element $a \in A$, we add an element $a' = a + m$ to S . For each element $b \in B$, we add an element $b' = b$ to S . And for each element $c \in C$, we add an element $c' = (-c - m)$ to S . We can see that if $(a+b+c)$ is equal to zero, then $(a' + b' + c')$ is also equal to zero. In order to construct a solution of PR2 from PR1, we need to show that whenever there are three elements in S that add up to zero, then the three elements came from each of the three different sets A , B and C .

Based on the way we have created a matching element in S for each element in A , B and C , we can see that a' lies between $0.5m$ to $1.5m$; b' lies between 0 to $0.5m$; and c' lies between $-1.5m$ to $-0.5m$. Now consider three elements x, y, z in S such that $x + y + z = 0$. Now, at the most, one of the elements in (x, y, z) could have been constructed from an element in A since, otherwise, the sum would be at least $2m - 1.5m = 0.5m$. Similarly, only one element could have been constructed from C , since otherwise the sum would be $-2m + 1.5m = -0.5m$. Also note that, out of (x, y, z) , one element must come from C , since elements constructed from A and B are always positive and we need at least one negative element to make the sum zero. This negative element can only come from C . Now, if we consider the possibility that one element in (x,y,z) is constructed from C and the other two are constructed from B , then again the sum of the two positive elements is at most m , whereas the element constructed from C is smaller than m , so the sum cannot

become zero. This leaves us with the only possibility that each of the elements in (x,y,z) was constructed from a different set. Hence, we have representatives constructed from all three sets A, B and C in (x,y,z) . Therefore, we have shown that if we can solve PR1 on the set S that we have constructed, we have a solution for PR2 also.

Now we can show a $O(N^2)$ algorithm for solving PR2 as follows: we sort the sets B and C. Consider an element $a \in A$. Now consider the set $B + a$ (the set of all numbers in B such that each number has a added to it). The set $(B + a)$ can be computed in $O(N)$ time. We now traverse the sets $(B + a)$ and C simultaneously to check if there is a common element between $(B + a)$ and C. This can be done in $O(N)$ time since these sets are sorted.

If we find a common element, we have found the triplet (a, b, c) such that $a + b = c$. Else we repeat the above procedure with the next element in A. Since we may end up repeating the above procedure for each of the elements in A, we end up with an overall complexity of $O(N^2)$ for our solution.

Now let us turn our attention to this month's topic of computational complexity and lower bounds for algorithms.

Computational complexity

Let us start with a simple children's game. You and your son are playing a number game. You tell your son that he can think of any positive number between zero and one million. You will ask him 20 questions to which he has to answer either "Yes" or "No" truthfully. You challenge your son that at the end of the 20 questions, you will be able to tell your son what number he had originally thought of. Your son is amazed at your ability to read his mind. Now how do you manage to answer correctly? The answer, of course, lies in the divide and conquer approach. With each question, you reduce the number range where the mystery number can possibly lie. For instance, your first question to your son is: "Is the number between 1 and 5,00,000?" An answer of "Yes" reduces your working range from 1 to 5,00,000, whereas an answer of "No" will reduce the answer range to 5,00,001 to one million. You know that you can always find the answer within your allotted 20 questions because one million is less than 2^{20} .

Okay, now that you have enthralled your son with your abilities, he asks you a question: Do you need 20 questions to arrive at the correct answer or can you do it with fewer questions? Now you are stumped. You know that the binary search algorithm that you have used has an upper bound of $O(\log N)$, and therefore you knew that 20 questions are sufficient to solve the problem. But your son's question to you now is whether 20 questions are necessary or can you do this with a fewer number of questions no matter what number he thinks up. This is essentially the realm of the computational complexity of problems. For this problem of finding a number in a given range, we can show that the computational complexity

has a lower bound of $O(\log N)$. That is, no matter what the algorithm you or anyone can else come up with, it will have a lower bound on the complexity as $O(\log N)$. If you had known this fact, you can tell your son that 20 questions are, in fact, necessary, for any arbitrary number he can think of. If you decide to challenge him to a context where only 19 questions are allowed, he can always win. I leave the detailed proof of this statement to interested readers since it would be fun to try this game with a friend/family member a few times first before trying to prove the statement rigorously.

Lower bounds of some well-known computational problems

Complexity analysis tries to prove that for a given problem, any algorithm capable of solving the problem correctly on all of its instances, must necessarily take a time that is greater than or equal to the established lower bounds for that problem. Basically, computational complexity theory establishes that there cannot exist an algorithm that can solve the problem in a time lower than the established lower bound, under the specified model of computation. For instance, it is widely known that any sorting algorithm based on comparison for sorting an array of N numbers has a lower bound of $O(N \log N)$. This means that there can exist no sorting algorithm based on the comparison of elements, which can sort faster than this.

There are a number of ways of establishing the lower bounds for any problem. Two of the well-known methods are the decision tree method and adversary argument. The decision tree method is well known and the interested reader can refer to any algorithm book for details on this. Let us take a look at the adversary argument for establishing lower bounds. Let us consider the well-known and simple problem of finding the maximum, given an array of N numbers.

Given below is the code snippet to find the maximum of a given set of N numbers, using only comparison between the numbers:

```
int find_maximum(int array[], int N)
{

    int max = A[0];

    for (int i = 0; i < N; i++)
    {
        if (a[i] > max)
            max = a[i];
    }
    return max;
}
```

It is evident that the algorithm given above uses exactly $(N-1)$ comparisons to arrive at the answer, since the *for* loop runs for $N-1$ iterations. Hence the algorithmic


complexity is $O(N)$. The interesting question for us is whether it is possible to find the maximum by doing fewer than $(N-1)$ comparisons. We can show the lower bound by using an adversary argument.

The idea behind an adversary argument is that we start off with an unspecified input and whenever the algorithm probes the input, the adversary gives an answer such that the algorithm is forced to work hard. The adversary must be consistent in its answers in the sense that there must exist at least one set of valid inputs on which the algorithm will exactly see the adversary answers matching with that input. The adversary's goal is to keep the algorithm uncertain of the answer for as long as possible. If the algorithm declares that it has found the answer without proceeding through the established lower bound number of probes, the adversary can always come up with a sequence of inputs that will invalidate the algorithm's answer.

For instance, consider a trivial example of a sequence of four numbers, namely $\{1, 2, 3, X\}$ and the algorithm needs to find the maximum. Assume that the algorithm performed two comparisons only, namely, between 1 and 2, and 2 and 3. It skipped the comparison of the current maximum 3 with the last element X in the array and declared the answer as 3. Since the algorithm has not probed $A[3]$, the adversary can fill it up with

any value for X , thereby defeating the algorithm. No matter how the algorithm tries, if it does not probe one of the elements of the array, the adversary can fill up that element with a value that will defeat the answer provided by the algorithm. Thus an adversary argument shows that any algorithm for finding the maximum in an array of elements has a lower bound of $O(N)$ comparisons for arriving at the correct answer.

For this month's takeaway problem, consider the well-known problem of deciding whether a given graph with N nodes is connected. The only question the algorithm can ask the adversary is of the form, "Does an edge exist between vertex u and vertex v ?" What is the best lower bound you can establish for this algorithm using an adversary argument?

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me, along with your feedback, to sandyasm_AT_yahoo_DOT_com. Till we meet again next month, happy programming! 

About the author:

Sandya Mannarswamy. The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.



Calling...

Share your knowledge

Inviting FOSS experts to write articles on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of www.openITis.com (aka linuxforu.com) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

lfyedit@efyindia.com

Few topics that top our list:

- ◆ Tips 'n' Tricks for software developets or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools



S.G. Ganesh

About the Java Overflow Bug

In this column, we'll discuss a common overflow bug in JDK, which surprisingly occurs in the widely used algorithms like binary search and mergesort in C-based languages.

How does one calculate the average of two integers, say i and j ? Trivial you would say: it is $(i + j) / 2$. Mathematically, that's correct, but it can overflow when i and j are either very large or very small when using fixed-width integers in C-based languages (like Java). Many other languages like Lisp and Python do not have this problem. Avoiding overflow when using fixed-width integers is important, and many subtle bugs occur because of this problem.

In his popular blog post [1], Joshua Bloch (Java expert and author of books on Java intricacies) writes about how a bug [2] in `binarySearch` and `mergeSort` algorithms was found in his code in `java.util.Arrays` class in JDK. It read as follows:

```

1: public static int binarySearch(int[] a, int key) {
2:     int low = 0;
3:     int high = a.length - 1;
4:
5:     while (low <= high) {
6:         int mid = (low + high) / 2;
7:         int midVal = a[mid];
8:
9:         if (midVal < key)
10:             low = mid + 1
11:         else if (midVal > key)
12:             high = mid - 1;
13:         else
14:             return mid; // key found
15:     }
16:     return -(low + 1); // key not found.
17: }

```

The bug is in line 6—"int mid = (low + high) / 2;". For large values of 'low' and 'high', the expression overflows and becomes a negative number (since 'low' and 'high' represent array indexes, they cannot be negative).

However, this bug is not really new—rather, it is usually not noticed. For example, the classic K & R book [3] on C has the same code (pg 52). For pointers, the expression $(low + mid) / 2$ is wrong and will result in compiler error, since it is not possible to add two pointers. So, the book's solution is to use subtraction (pg 113):

```
mid = low + (high-low) / 2
```

This finds 'mid' when 'high' and 'low' are of the same sign

(they are pointers, they can never be negative). This is also a solution for the overflow problem we discussed on Java.

Is there any other way to fix the problem? If 'low' and 'high' are converted to unsigned values and then divided by 2, it will not overflow, as in:


```
int mid = ( (unsigned int) low + (unsigned int) high) / 2;
```

But Java does not support unsigned numbers. Still, Java has an unsigned right shift operator (`>>>`)—it fills the right-most shifted bits with 0 (positive values remain as positive numbers; also known as 'value preserving'). For the Java right shift operator `>>`, the sign of the filled bit is the value of the sign bit (negative values remain negative and positive values remain positive; also known as 'sign-preserving'). Just as an aside for C/C++ programmers: C/C++ has only the `>>` operator and it can be sign or value preserving, depending on implementation. So we can use the `>>>` operator in Java:

```
int mid = (low + high) >>> 1;
```

The result of $(low + high)$, when treated as unsigned values and right-shifted by 1, does not overflow!

Interestingly, there is another nice 'trick' to finding the average of two numbers: $(i \& j) + (i \wedge j) / 2$. This expression looks strange, doesn't it? How do we get this expression? Hint: It is based on a well-known Boolean equality, for example, as noted in [4]: " $(A \text{ AND } B) + (A \text{ OR } B) = A + B = (A \text{ XOR } B) + 2 (A \text{ AND } B)$ ".

A related question: How do you detect overflow when adding two *ints*? It's a very interesting topic and is the subject for next month's column. 

References:

- googleresearch.blogspot.com/2006/06/extra-extra-read-all-about-it-nearly.html
- bugs.sun.com/bugdatabase/view_bug.do?bug_id=5045582
- The C Programming Language, Brian W. Kernighan, Dennis M. Ritchie, Prentice-Hall, 1988.
- home.pipeline.com/~hbaker1/hakmem/boolean.html#item23

About the author:

S G Ganesh is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill in December 2007. You can reach him at sgganesh@gmail.com.

A Voyage to the Kernel



Part 9

Segment – 2.2, Day 8

In the last column, we had discussed some basic algorithms and methodologies. Now we will generalise the scheme of an algorithm. It will have the following properties:

- Arranged in an ordered sequence that you can number (as steps)
- Unambiguous and well-defined
- Halts in finite time (i.e., the algorithm terminates!)

And we will generally stumble upon the following algorithmic operations:

- Sequential operations—where instruction sets are executed in order
- Conditional operations—that ask a true/false question and then select the next instruction based on the answer
- Iterative operations (loops)—that repeat the execution of a set of instructions

It should be emphasised that, as we discussed before, not every problem has a ‘good’ algorithmic solution. There are:

- Halting problems (unsolvable problems)—for which no algorithm exists to solve the problem.
- Travelling salesman problems (intractable problems)—algorithms that take too long to solve the problem.
- Problems with no known algorithmic solution.

Pseudo code structure

The algorithm may contain the following elements: IF-THEN-ELSE

Binary choice on a given Boolean condition

is indicated by: IF, THEN, ELSE, and ENDIF. The general representation is:

IF condition THEN

sequence 1

ELSE

sequence 2

ENDIF

Please note that the ELSE keyword and “sequence 2” are optional.

WHILE: The loop is executed only if the condition is true. The ‘WHILE construct’ is used to specify a loop with a test at the top. WHILE and ENDWHILE are to identify the beginning and ending of the loop. The general form for WHILE is:

WHILE condition

sequence

ENDWHILE

CASE: In order to handle mutually exclusive conditions, we can go for CASE. CASE, OF, OTHERS, and ENDCASE are the keywords commonly used for this. The alternatives available will be represented with the help of ‘conditions’.

CASE expression OF

```

condition 1 : sequence 1
condition 2 : sequence 2
...
condition n : sequence n
OTHERS:
default sequence

```

ENDCASE

REPEAT-UNTIL: This is quite akin to WHILE, but here the operation is performed at the bottom of the loop.

REPEAT

```
sequence
```

UNTIL condition

FOR: We employ FOR and ENDFOR in our algorithm for iterating the code block for a specific number of times. It is also called the “counting” loop.

FOR iteration bounds

```
sequence
```

ENDFOR

EXCEPTION HANDLING: The following code elucidates this.

```

BEGIN
    statements
EXCEPTION
    WHEN exception type
        statements to handle exception
    WHEN another exception type
        statements to handle exception
END

```

You can also utilise methods like NESTED CONSTRUCTS, INVOKING SUBPROCEDURES to construct elegant and powerful algorithms.

We can define seven golden rules for writing algorithms:

- Use good code and good English.
- Ignore unnecessary details
- Take advantage of programming shorthand
- Be context specific
- Don't lose sight of the base model
- Always check for balance (can be easily implemented in a language)

Implementation

Let's write the code that produces a random number. Random number generation is an important tool in




Figure 1: The free Pascal compiler

simulation building and scientific computing. In order to elucidate the implementation of the algorithm, I have written the code in Pascal.

```

program random_no (input, output);
const m=100000000; m1=10000; q=31415821;
var i, p, No: integer;
function mult(r, s: integer): integer;
begin
    var r1, r0, s1, s0: integer;
begin
    r1:=r div m1; r0:=r mod m1;
    s1:=s div m1; s0:=s mod m1;
    mult:=( ((r0*s1+r1*s0) mod m1)*m1+r0*s0)
end;
function random : integer;
begin
    p:=(mult(p, q)+1) mod m;
    random :=p;
end;
begin
    read(No, p);
    for i:=1 to No do writeln(random)
end.

```

It's assumed that you are skilled in grasping the meaning of simple codes. You can try this in your preferred Pascal compiler. (Figure 1 shows the code in Free Pascal Compiler, which I installed for trial.)

Euclid's Algorithm: A response to the queries

Many readers who tried the algorithms in the reference books (suggested in the last column) posed the question concerning the simplification of Euclid's algorithm and the base problem. As it was given as a problem, many of 'our passengers' got wedged between the lines!

For those who have not tried the exercises (if you don't have the book, you can find many online resources), here is a prologue to the problem:

Let a and b be integers, with $a \geq b \geq 0$. Using the

division with the remainder property, define the integers $r_0, r_1, \dots, r_{\lambda+1}$, and q_1, \dots, q_λ , where $\lambda \geq 0$, as follows:

```
a = r0,
b = r1,
r0 = r1 q1 + r2 (0 < r2 < r1),
.
.
.
ri-1 = ri qi + ri+1 (0 < ri+1 < ri),
.
.
.
rλ-2 = rλ-1 qλ-1 + rλ (0 < rλ < rλ-1),
rλ-1 = rλ qλ (rλ+1 = 0).
```

From the given definitions $\lambda = 0$ if $b = 0$, and $\lambda > 0$ for all other cases. Then we have $r_\lambda = \gcd(a, b)$. Moreover, if $b > 0$, then $\lambda \leq \log b / \log \phi + 1$, where $\phi := (1 + 5)/2 \approx 1.62$.

Algorithm: If we could furnish input as a and b , where a and b are integers such that $a \geq b \geq 0$, we can compute $d = \gcd(a, b)$ as exemplified below:

```
r ← a, r ← b
while r = 0 do
    r ← r mod r
    (r, r) ← (r, r)
d ← r
output d
```

Designing algorithms for specific problems

We have the Legendre differential equation as:

$$\frac{d}{dx} \left[(1-x^2) \frac{d}{dx} P_n(x) \right] + n(n+1) P_n(x) = 0$$

Also, please note that the Legendre equation may be represented by:

$$ax^2 + by^2 + cz^2 = 0$$

And we also have the Sturm-Liouville equation that is a real second-order linear equation of the form:

$$-\frac{d}{dx} \left[P(x) \frac{dy}{dx} \right] + q(x)y = \lambda w(x)y$$

Here, y is a function of the free variable x , and $p(x)$, $q(x)$ and $w(x)$ are the functions of x . We will be

handling the Sturm-Liouville problem that intends to find the value of λ . Please refer to mathworld.wolfram.com/Sturm-LiouvilleEquation.html for more information. Now our concern is to build an algorithm. Let's do this in C.

Problem: To solve the Legendre equation with the simplest algorithm for the Sturm-Liouville equation,

/* Solving the Legendre equation with the simplest algorithm for the Sturm-Liouville equation

Code written for A Voyage to Kernel - 8*/

```
#include <stdio.h>
#include <math.h>
#define NOMAX 1000

main()
{
    int i,no,inc;
    double dl=1e-5;
    double u[NOMAX];
    double z,r,s,uu,t,f0,f1;
    void salgo();

    /* Initialization of our problem */

    no = NOMAX;
    z = 3.0/(no-1);
    s = 1.5;
    r = 0.5;
    t = 0.5;
    inc = 0;
    u[0] = -1;
    u[1] = -1+z;
    uu = r;
    salgo (no,z,uu,u);
    f0 = u[no-1]-1;
    while ((fabs(t) > dl)
    {
        uu = (r+s)/2;
        salgo (no,z,uu,u);
        f1 = u[no-1]-1;
        if ((f0*f1) < 0)
        {
            s = uu;
            t = s-r;
        }
        else
        {
            r = uu;
            t = s-r;
            f0 = f1;
        }
        inc = inc+1;
    }
    printf("%4d %16.8lf %16.8lf %16.8lf\n", inc,uu,t,f1);
```



```

}
void salgo (no,z,uu,u)

/* Consider simplest algorithm for the Sturm-Liouville equation.*/

int no;
double z,uu;
double u[];
{
int i;
double z2,q,p,p1,x;
q = uu*(1+uu);
z2 = 2*z*z;
for (i = 1; i <no-1; ++i)
{
p1 = -2*x*z;
p = 2*(1-x*x);
x = -1+i*z;
u[i+1] = ((2*p-z2*q)*u[i]+(p1-p)*u[i-1])/(p1+p);
}
}

```

You can extend this algorithmic method to solve other problems as many equations like the Bessel equation, which is given by...

$$x^2 y'' + xy' + (\lambda^2 x^2 - \nu^2) y = 0$$

...can be represented in the Sturm-Liouville form as

$$(xy') + (\lambda^2 x^2 - \nu^2/x)y = 0$$

We will now try out Lagrange interpolation using the Aitken method. I'm sure you must have done a simple form of this already. But this one is slightly diverse.

```

/* Main program for the Lagrange interpolation with the Aitken method.*/

#include <stdio.h>
#include <math.h>
#define NOMAX 3

main()

{
int i;
double f, dfi;
double h = 0.5, x = 0.9;
double xi[NOMAX], fi[] = {1.1, 0.9, 0.7, 0.5, 0.3};
void aitken_method_kernel_voyage();
for (i = 0; i < NOMAX; i++)
{
xi[i] = h*i;
}
aitken_method_kernel_voyage(NOMAX, xi, fi, x, &f, &dfi);

```

To my readers

Thanks for your response, feedback and suggestions. But I am extremely sorry that I am not in a position to reply to all the queries made through e-mails. The major problem is that e-mails often get messed up. Hence, I may miss your e-mail. Sometimes I mark your mails, but the new e-mails often mask them. To avoid this problem, I have created a new platform for answering questions:

aasisvinayak.com/new_zone/forum.php

Considering the suggestions from readers, the next column will focus on notations and numerical analysis. But the emphasis will be on the designing of algorithms, as this is a kernel-programming column and not a section that addresses pure scientific computing.

```

printf("%10.6lf %10.6lf %10.6lf\n", x, f, dfi);
}
void aitken_method_kernel_voyage(n, xi, fi, x, f, dfi)
int n;
double xi[], fi[];
double x;
double *f, *dfi;
{
int i, j;
double fn[NOMAX];
double x1, x2, f1, f2;
for (i = 0; i <= (n-1); ++i)
{
fn[i] = fi[i];
}
for (i = 0; i <= (n-2); ++i)
{
for (j = 0; j <= (n-i-2); ++j)
{
x1 = xi[j];
x2 = xi[j+i+1];
f1 = fn[j];
f2 = fn[j+1];
fn[j] = (x-x1)/(x2-x1)*f2+(x-x2)/(x1-x2)*f1;
}
}
*f = fn[0];
*dfi = (fabs(*f-f1)+fabs(*f-f2))/2;
}

```

You can see the output by executing this code in your compiler. (You can try changing the values assigned and see what happens.)

Iterations in coding

As we discussed before, iteration plays an important role when it comes to designing algorithms. Now, we will try to find the value of Pi using this.

```

/* Estimating Pi - Iteration*/

```




Figure 2: A SAMPLE TREE

```

#include <stdlib.h>
#include <stdio.h>
#include <math.h>
#include <string.h>
#define STD 23742246
main(int argc, char* argv)
{
    int no=0;
    double x,y;
    int i,count=0;
    double z;
    double pi;
    printf("Enter the number of iterations to be done to estimate the value of
pi: ");
    scanf("%d",&no);
    srand(STD);
    count=0;
    for (i=0; i<no; i++) {
        x = (double)rand()/RAND_MAX;
        y = (double)rand()/RAND_MAX;
        z = x*x+y*y;
        if (z<=1) count++;
    }
    pi=(double)count/no*4;
    printf("No of trials= %d , The Estimated value of pi is %g \n",no,pi);
}
    
```

You can see that the estimated value of Pi approaches the correct value when the number of steps in the procedure increases (dynamically). If you look at the code carefully, you will see that the code terminates after a finite period (provided the input is finite!):

```

aasisvinayak@GNU-BOX:~/Documents/Desktop$ ./pi
Enter the number of iterations to be done to estimate the value of pi: 44
No of trials= 44 , The Estimated value of pi is 2.81818
aasisvinayak@GNU-BOX:~/Documents/Desktop$ ./pi
Enter the number of iterations to be done to estimate the value of pi: 1234567
No of trials= 1234567 , The Estimated value of pi is 3.14508
    
```

We have seen the implemented codes. Now, try writing the algorithms behind the codes. As you have the code, it is quite easy!

Another way of visualising our algorithms is the Tree Model. Beginners can opt for this since this will never cause any bewilderment. Figure 2 shows a model of A

SAMPLE TREE. This will lend a hand to such users to comprehend the idea of branching.

Today's problem

We have a theorem: If Ω be the set of all exact execution paths for a defined 'A' on input x, then...

$$\sum_{\lambda \in \Omega} 2^{-|\lambda|} \leq 1$$

Can you prove this? (Note that all symbols have the usual meaning that we used for other algorithms/problems.)

Clues:

Start with...

$$\alpha_i := \sum_{\lambda \in \Omega} 2^{-|\lambda|} \leq 1$$

...and deduce

$$\alpha_i \leq 1$$


Then, try to get...

$$\sum_{\lambda \in \Omega} 2^{-|\lambda|} = \lim_{i \rightarrow \infty} \alpha_i \leq 1$$

...by considering ...

$$\alpha_i = 2^{-i} |S_k|$$

If you can reach this juncture, the rest will follow. For all these, you can assume that A halts with probability α on input x. If $\alpha = 1$, we define the distribution $P: \Omega \rightarrow [0, 1]$.

Beginners need not worry much. We will be dealing with complex problems much later in our Voyage! 

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

FOSS Yellow Pages


The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



TARIFF

Category Listing

ONE Category	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category	Rs 1,000

Value-add Options

LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

✂ Tear & Send

ORDER FORM

Tear & Send ✂

Organisation Name (70 characters): _____

Description (30 words): _____

Email: _____ Website: _____

STD Code: _____ Phone: _____ Mobile: _____

Address (will not be published): _____

City/Town: _____ Pin-code: _____

CATEGORIES

- | | | |
|---|--|---|
| <input type="checkbox"/> CONSULTANTS | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING | <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| <input type="checkbox"/> CONSULTANT (FIRM) | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS | <input type="checkbox"/> LINUX-BASED WEB-HOSTING | <input type="checkbox"/> TRAINING FOR CORPORATE |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS | <input type="checkbox"/> THIN CLIENT SOLUTIONS |

Please find enclosed a sum of Rs. _____ by DD/ MO//crossed cheque* bearing the No. _____ dt. _____ in favour of

EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs. 50 on non-metro cheque) towards the cost of _____ FOSS Yellow Pages advertisement(s)

or charge my credit card ☐ VISA ☐ Master Card Please charge Rs. _____

against my credit card No. _____

C V V No. _____ (Mandatory)

Date of Birth ____ / ____ / ____ (dd/mm/yy) Card Expiry Date ____ / ____ (mm/yy)

Signature (as on the card)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

To Book Your Listing, Call: Dhiraj (Delhi: 09811206582), Somaiah (B'lore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Consultant (Firm)

OS3 Infotech

•Silver Solutions Partner for Novell
•High Availability Computing
Solutions •End-to-end Open Source
Solutions Provider •Certified Red
Hat Training Partner •Corporate and
Institutional Training

Navi Mumbai

Mobile: 09324113579

Email: info@os3infotech.com

Web: www.os3infotech.com

Taashee Linux Services

100% Support on LINUX ,OSS
& JBOSS related projects. We
specialize in high-availability and
high-performance clusters,remote
and onsite system management,
maintenance services,systems
planning, Linux & JBOSS consulting
& Support services.

Hyderabad

Mobile: 09392493753, Fax: 040-40131726

Email: manojkummar@taashee.com

Web: www.taashee.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer comptek
wibrain B1 UMPC with Linux,Touch
Screen, 1 GB RAM 60GB, Wi-Fi,
Webcam, upto 6 hour battery
(opt.), USB Port, max 1600x1200
resolution, screen 4.8", 7.5"x3.25"
Size, weight 526 gm.

New Delhi

Mobile: 09968756177, Fax: 011-26187551

Email: comptekdelhi@compteki.com

Web: www.compteki.com

or www.compteki.in

Education & Training

Aptech Limited

IT, Multimedia and Animation
Education and Training

Mumbai

Tel: 022-28272300, 66462300

Fax: 022-28272399

Email: customer@aptech.ac.in

Web: www.aptech-education.com,

www.arena-multimedia.com

IT-Campus : Academy of Information Technology

IT training and solution company with
over 12 years of experience. - RHCE

•Software Training •Hardware
Training •Multimedia And Animation

•Web Designing •Financial
Accounting

Kota (Raj.)

Tel: 0744-2503155, Mobile: 09828503155

Fax: 0744-2505105

Email: m_trilok@yahoo.com

Web: www.doeacc4u.com

Mahan Computer Services (I) Limited

Established in 1990, the organization
is primarily engaged in Education
and Training through its own &
Franchise centres in the areas of IT
Software, Hardware, Networking,
Retail Management and English. The
institute also provides customized
training for corporates.

New Delhi

Tel: 011-25916832-33

Email: info@mahanindia.com

Web: www.mahanindia.com

Enterprise Communication Solutions

Aware Consultants

We specialize in building and
managing Ubuntu/Debian
Linux servers and provide good
dependable system administration.
We install and maintain in-house
corporate servers. We also provide
dedicated and shared hosting as well
as reliable wireless/hybrid networking.

Bangalore

Tel: 080-26724324

Email: sales@aware.co.in

Web: www.aware.co.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty
at the Indian Institute of Science,
Bangalore and carry over eight
decades of experience and
fundamental knowledge in the field
of DSP and Telecommunication.
ESQUBE plays a dominant role in
the creation of IP in the domain of
Sensors, Signals and Systems.

Bangalore

Tel: 080-23517063

Email: info@esqube.com

Web: www.esqube.com

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open
source Enterprise Groupware
Mobile push, Email Syncing of
Contacts/Calendar/Tasks with
mobiles

•Mail Archival •Mail Auditing •Instant
Messaging

New Delhi

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

netCORE The Innovation Company

Netcore Solutions Pvt Ltd

No.1 company for providing Linux
Based Enterprise Mailing solution
with around 1500+ Customer all over
India. Key Solutions:

•Enterprise Mailing and Collaboration
Solution •Hosted Email Security •Mail
Archiving Solution •Push Mail on
Mobile •Clustering Solution

Mumbai

Tel: 022-66628000

Mobile: 09322985222

Email: kalpit@netcore.co.in

Web: www.netcore.co.in



Red Hat India Pvt Ltd

Red Hat is the world's leading
open source solutions provider.
Red Hat provides high-quality,
affordable technology with its
operating system platform, Red
Hat Enterprise Linux, together with
applications, management and
Services Oriented Architecture (SOA)
solutions, including JBoss Enterprise
Middleware. Red Hat also offers
support, training and consulting
services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis TechnoLab is the first of its
kind, state-of-the-art infrastructure,
Hardware, Networking and
I.T Security training institution
headquartered in Kolkata. TechnoLab

is the training division of Xenitis group
of Companies. It is the proud owner
of 'Aamar PC', the most popular
Desktop brand of Eastern India.
These ranges of PC's are sold in
the west under the brand name of
'Aamchi PC', in the north as 'Apna
PC' and in the south as 'Namma PC'.

Kolkata

Tel: 033-22893280

Email: srinku@xenitisgroup.com

Web: www.technolabindia.com

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd

Open Source Solutions Provider.
Red Hat Ready Business Partner.
Mail Servers/Anti-spam/GUI
interface/Encryption, Clustering &
Load Balancing - SAP/Oracle/Web/
Thin Clients, Network and Host
Monitoring, Security Consulting,
Solutions, Staffing and Support.

New Delhi

Tel: +91-11-26494549

Fax: +91-11-4175 1823

Mobile: +91-9873839960

Email: sales@aisplglobal.com

Web: www.aisplglobal.com

Advent Infotech Pvt Ltd

Advent has an experienced techno-
marketing team with several years of
experience in Networking & Telecom
business, and is already making
difference in market place. ADVENT
qualifies more as Value Added
Networking Solution Company, we
offers much to customers than just
Routers, Switches, VOIP, Network
Management Software, Wireless
Solutions, Media Conversion, etc.

New Delhi

Tel: 46760000, 09311166412

Fax: 011-46760050

Email: marketingsupport@adventelectronics.com

Web: www.adventelectronics.com

Asset Infotech Ltd

We are an IT solution and training
company with an experience of 14
years, we are ISO 9001: 2000. We
are partners for RedHat, Microsoft,
Oracle and all Major software
companies. We expertise in legal
software ans solutions.

Dehradun

Tel: 0135-2715965, Mobile: 09412052104

Email: piyush@asset.net.in

Web: www.asset.net.in

To advertise in this section, please contact
Somaiah (Bangalore) **09986075717**
Dhiraj (Delhi) **09811206582**

FOSS Yellow Pages

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi

Tel: 011-42235156

Email: ashish.gupta@bakbone.com

Web: www.bakbone.com

Clover Infotech Private Limited

Clover Infotech is a leading technology services and solutions provider. Our expertise lies in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Services and Application Development Services.

Mumbai

Tel: 022-2287 0659, Fax: 022-2288 1318

Mobile: +91 99306 48405

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

Duckback Information Systems Pvt Ltd

A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix, Adobe, Redhat, Novell, Symantec, McAfee, Computer Associates, Veritas, Sonic Wall

Kolkata

Tel: 033-22835069, 9830048632

Fax: 033-22906152

Email: asis@duckback.com

Web: www.duckback.co.in

FOSS Yellow Pages

HBS System Pvt Ltd

System Integrators & Service Provider. Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, Trend Micro, Symantec Partners of SUN for their new startup E-commerce initiative Solution Provider on REDHAT, SOLARIS & JAVA

New Delhi

Tel: 011-25767117, 25826801/02/03

Fax: 25861428

Email: amittal@hbsindia.com



Ingres Corporation

Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are hallmarks of an open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10,000 customers in the United States and internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66459537

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

Keen & Able Computers Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: info@keenable.com

Web: www.keenable.com

LDS Infotech Pvt Ltd

Is the authorised partner for RedHat Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Corel etc. Software Services Offered: •Collaborative Solutions •Network Architecture •Security Solutions •Disaster Recovery •Software Licensing •Antivirus Solutions.

Mumbai

Tel: 022-26849192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

Pacer Automation Pvt Ltd

Pacer is leading providers of IT Infrastructure Solutions. We are partners of HP, Redhat, Cisco, Vmware, Microsoft and Symantec. Our core expertise exists in, Consulting, building and Maintaining the Complete IT Infrastructure.

Bangalore

Tel: 080-42823000, Fax: 080-42823003

Email: solutions@pacerautomation.com

Web: www.pacerautomation.com



Red Hat India Pvt Ltd

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai

Tel: 022-39878888

Email: marketing-in@redhat.com

Web: www.redhat.in

Srijan Technologies Pvt Ltd

Srijan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.

New Delhi

Tel: 011-26225926, Fax: 011-41608543

Email: business@srijan.in

Web: www.srijan.in

TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai

Tel: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions caters of all kind of industry vehicles.

New Delhi

Tel: 011-46571313, Fax: 011-41620171

Email: sales@tetra.in

Web: www.tetra.in

Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212098104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

FOSS Yellow Pages

Want to register your organisation in
FOSS Yellow Pages For **FREE***

Call: Dhiraj (Delhi) 09811206582 Somaiah (Bangalore) 09986075717
or mail: dhiraj.khare@efyindia.com, somaiah.km@efyindia.com

**Offer for limited period.*

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) **09811206582**, Somaiah (Bangalore) **09986075717**

Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

Chennai

Tel: 044-42210000, Fax: 28144986

Email: info@veeras.com

Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting

ManasHosting is a Bangalore-based company that is dedicated in helping small and midsize business companies to reach customers online. We believe that by creating a website, all you have is just web presence; but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs

Bangalore

Tel: 080-42400300

Email: enquiry@manashosting.com

Web: www.manashosting.com

Linux Desktop

Indserve Infotech Pvt Ltd

OpenLx Linux with Kalcuate (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500 + tax as special scheme), Friendly (Graphical UserInterface) and Secure (Virus free).

New Delhi

Tel: 011-26014670-71, Fax: 26014672

Email: info@openlx.com

Web: www.openlx.com

Linux Experts

Intaglio Solutions

We are the training and testing partners of RedHat and the first to conduct RHCCS exam in delhi for the first time ever.

New Delhi

Tel: 011-41582917, 45515795

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises

Authorized distributors for Red Hat and JBoss range of products. We also represent various OS's Applications and Developer Tools like SUSE, VMWare, Nokia Qt, MySQL, Codeweavers, Ingres, Sybase, Zimbra, Zend-A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Pathscale/Qlogic and Intel Compilers, Scalix-Messaging solution on Linux Platform.

Bangalore

Mobile: +91-9845009939, +91-9343861758

Email: sales@gte-india.com

Web: www.gte-india.com

Taurusoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details

Mumbai

Mobile: 09869459928, 09892697824

Email: taurusoft@gmail.com

Web: www.taurusoft.netfirms.com

Software Subscriptions

Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, JBoss, Oracle, ARCserve Backup, AntiVirus for Linux, Verisign/Thawte/GeoTrust SSL Certificates and many other original software licenses.

Mumbai

Tel: 022-25001812, Mobile: 09821097238

Email: bluechip@vsnl.com

Web: www.bluechip-india.com

Software Development

Carizen Software (P) Ltd

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

Chennai

Tel: 044-24958222, 8228, 9296

Email: info@carizen.com

Web: www.carizen.com

DeepRoot Linux Pvt Ltd

DeepRoot Linux is a seven year old GNU/Linux and Free Software company based in Bangalore. We develop Free Software products that are quick-to-deploy and easy-to-use.

Bangalore

Tel: 080-40890000

Email: shivanand@deeproot.in

Web: www.deeproot.in



InfoAxon Technologies Ltd

InfoAxon designs, develops and supports enterprise solutions stacks leveraging open standards and open source technologies. InfoAxon's focus areas are Business Intelligence, CRM, Content & Knowledge Management and e-Learning.

Noida

Tel: 0120-4350040, Mobile: 09810425760

Email: sales@infoaxon.com

Web: <http://opensource.infoaxon.com>

Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.

Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integramicro.com

Web: www.integramicroservices.com

iwebtune.com Pvt Ltd

iwebtune.com is your one-stop, total web site support organisation. We provide high-quality website services and web based software support to any kind of websites, irrespective of the domain or the industry segments.

Bangalore

Tel: 080-4115 2929

Email: santosh@iwebtune.com

Web: www.iwebtune.com

Want to register your organisation in
FOSS Yellow Pages
For **FREE**
CALL: DHIRAJ (DELHI) 09811206582
SOMAIAH (BANGALORE) 09986075717
*Offer for limited period.

Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national sales & support partner for BitDefender Antivirus products.

New Delhi

Tel: 011-26288583, Fax: 011-26219396

Email: isales@unistal.com

Web: www.unistal.com

Software and Web Development

Bean eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

Navi Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: infodesk@beanarchitect.com

Web: www.beanarchitect.com

CATEGORIES FOR FOSS YELLOW PAGES

- ☐ Consultants
- ☐ Consultant (Firm)
- ☐ Embedded Solutions
- ☐ Enterprise Communication Solutions
- ☐ High Performance Computing
- ☐ IT Infrastructure Solutions
- ☐ Linux-based Web-hosting
- ☐ Mobile Solutions
- ☐ Software Development
- ☐ Training for Professionals
- ☐ Training for Corporate
- ☐ Thin Client Solutions

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour; we also have a Customer Care Center in India for any kind of after sales help. We are already selling it world over with over 65,000 copies sold. It comes with FREE Domain Name, Web Hosting and Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

Patiala

Mobile: 91-9780531682

Email: pardeep@mrsite.co.in

Web: www.mrsite.co.in

Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

New Delhi

Tel: 011-41648668, 66091565

Email: ceo@salahsoftware.com

Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves

The 'System Integration' business unit offers end-to-end Solutions on Desktops, Servers, Workstations, HPC Clusters, Render Farms, Networking, Security/Surveillance & Enterprise Storage. With our own POWER-X branded range of Products, we offer complete Solutions for Animation, HPC Clusters, Storage & Thin-Client Computing.

Mobile: 09880715253

Email: ranga@digitalwaves.in

Web: www.digitalwaves.in

Enjoy Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment.

Active Dealer Channel all over India.

Gujarat

Tel.: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: info@enjoyworld.com

Web: www.enjoyworld.com

Training for Corporate

Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also providing training to the individuals and corporate on other open source technologies like PHP, MySQL etc.

Ahmedabad

Tel: 079-27545455—66

Fax: 079-27545488

Email: info@bascombridge.com

Web: www.bascombridge.com

Brainnet

Kolkata

Tel: 033-40076450

Email: brainnet@brainware-india.com

Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTNL

MTNL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom; Management personnel. CETTM has AC lecture halls, computer Labs and residential facility.

Mumbai

Tel: 022-25714500, 25714586, 25714585, 25714586

Fax: 022-25706700

Email: contact@cettm.mtnl.in

Web: http://cettm.mtnl.in/infra

Complete Open Source Solutions

RHCT, RHCE and RHCSS training.

Hyderabad

Tel: 040-66773365, 9849742065

Email: nayak.sujeet@gmail.com

Web: www.cossindia.com

ElectroMech

Redhat Linux and open source solution, RHCE, RHCSS training and exam center, Ahmedabad and Vadodara

Ahmedabad

Tel: 079-40027898

Email: electromech@electromech.info

Web: www.electromech.info

Focuz Infotech

Focuz Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology Training, Certifications and live projects to students and corporates, since 2000.

Cochin

Tel: 0484-2335324

Email: enquiry@focuzinfotech.com

Web: www.focuzinfotech.com

Gujarat Infotech Ltd

GIL is a IT company and 17 years of experience in computer training field. We have experience and certified faculty for the open Source courses like Redhat, Ubuntu, and PHP, MySQL

Ahmedabad

Tel: 079-27452276, Fax: 27414250

Email: info@gujaratinfotech.com

Web: www.gujaratinfotech.com

Lynus Academy Pvt Ltd

India's premier Linux and OSS training institute.

Chennai

Tel: 044-42171278, 9840880558

Email: contactus@lynusacademy.com

Web: www.lynusacademy.com

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore

Tel: 080-22428538, 26600839

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd

Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality deliverables in the fields of BPO, Software and Networking,

while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the spheres of Information Technology (IT) and Information Technology Enabled Services (ITES).

Chennai

Tel: 044-45582525

Email: info@mazenetsolution.com

Web: www.mazenetsolution.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore

Tel: 080-41146565, 32719516

Email: info@netwebindia.com

Web: www.netwebindia.com

New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global foot print engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning, Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

Network NUTS

India's only Networking Institute by Corporate Trainers. Providing



FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Corporate and Open classes for RHCE / RHCSS training and certification. Conducted 250+ Red Hat exams with 95% result in last 9 months. The BEST in APAC.

New Delhi

Tel: 46526980-2

Mobile: 09310024503, 09312411592

Email: info@networknuts.net

Web: www.networknuts.net

STG International Ltd

An IT Training and Solution Company. Over an experience of 14 years. We are ISO 9001:2000 Certified. Authorised Training Partners of Red Hat & IBM-CEIS. We cover all Software Trainings.

New Delhi

Tel: 011-40560941-42, Mobile:

09873108801

Email: rakhi@stg.in

Web: www.stgonline.com

www.stgglobal.com

TNS Institute of Information Technology Pvt Ltd

Join RedHat training and get 100% job guarantee. World's most respected Linux certification. After RedHat training, you are ready to join as a Linux Administrator or Network Engineer.

New Delhi

Tel: 011-3085100, Fax: 30851103

Email: nks@tiit.co.in

Web: www.tiit.co.in

Webel Informatics Ltd

Webel Informatics Ltd (WIL), a Government of West Bengal Undertaking. WIL is Red Hat Training Partner and CISCO Regional Networking Academy. WIL conducts RHCE, RHCSS, CCNA, Hardware and Software courses.

Kolkata

Tel: 033-22833568, Mobile: 09433111110

Email: enquiry@webelinformatics.com

Web: www.webelinformatics.com

To advertise in this section,
please contact

Somaiah (Bangalore)

09986075717

Dhiraj (Delhi)

09811206582

Training for Professionals

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

Dehradun

Tel: 0135-2673712, Mobile: 09760099050

Web: www.agamtecindia.com

Amritha Institute of Computer Technology

Amritha Technologies provides an extensive training in high end certification programs and Networking Solutions like Redhat Linux, Redhat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and so on with a strong focus on quality standards and proven technology processes with most profound principles of Love and Selfless Service.

Mobile: 09393733174

Email: aict.hybd@amrita.ac.in

Web: www.amritahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to the need for industries as well as Students for trainings on IT, CISCO certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handicraft.

Bhopal

Tel: 0755-2661412, 2661559

Fax: 0755-4220022

Email: crisp@crispindia.com

Web: www.crispindia.com

Center for Open Source Development And Research

Linux, open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective business ready solutions caters of all kind of industry verticals.

New Delhi

Mobile: 09312506496

Email: info@cfosdr.com

Web: www.cfosdr.com

Cisconet Infotech (P) Ltd

Authorised Red Hat Study cum Exam Centre. Courses Offered: RHCE,

RHCSS, CCNA, MCSE

Kolkata

Tel: 033-25395508, Mobile: 09831705913

Email: info@cisconetinfo.com

Web: www.cisconetinfo.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state Of The Art IT Infrastructure Flexible Batch Timings Available.. Leading Networking Institute in Marathwada

Aurangabad

Tel: 0240-3299509, 6621775

Email: aurangabad@cmsinstitute.co.in

Web: www.cmsaurangabad.com

Cyber Max Technologies

OSS Solution Provider, Red Hat Training Partners, Oracle, Web, Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs

Bikaner

Tel: 0151-2202105, Mobile: 09928173269

Email: cmtech.bikaner@gmail.com,

kr.gupta.ashish@gmail.com

Disha Institute

A franchisee of Unisoft Technologies, Providing IT Training & Computer Hardware & Networking

Dehradun

Tel: 3208054, 09897168902

Email: thedishainstitute@gmail.com

Web: www.unisofttechnologies.com

EON Infotech Limited (TECHNOSchool)

TechnoSchool is the most happening Training Centre for Red Hat (Linux- Open Source) in the Northern Region. We are fully aware of the Industry's requirement as our Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

Chandigarh

Tel: 0172-5067566-67, 2609849

Fax: 0172-2615465

Email: info@technoschool.net

Web: http://technoschool.net

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer Hardware Networking, Mobile Phone Maintenance & International Certifications

Jaipur

Tel: 0141-3213378

Email: franchise_gt@gteducation.net

Web: www.gteducation.net

HCL Career Development Centre Bhopal

As the fountainhead of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a Leader can transform you into a Leader". HCL CDC is a formalization of this experience and credo which has been perfected over three decades.

Bhopal

Tel: 0755-4094852

Email: bhopal@hclcdc.in

Web: www.hclcdc.in

IINZTRIX E Technologies Pvt Ltd

No. 1 Training provider in this region.

meerut

Tel: 0121-4020111, 4020222

Mobile: 09927666664

Email: jai@iintrix.com

Web: www.iintrix.com

Indian Institute of Job Oriented Training Centre

Ahmedabad

Tel: 079-40072244—2255—2266

Mobile: 09898749595

Email: info@iijit.net

Web: www.iijit.net

Institute of Advance Network Technology (IANT)

•Hardware Engg. •Networking
•Software Engg. •Multimedia
Training.

Ahmedabad

Tel: 079-32516577, 26607739

Fax: 079-26607739

Email: contact@iantindia.com

Web: www.iantindia.com

IPCC

Bridging Gap with professionals.

Lucknow

Tel: 0522-3919496

Email: ipccko@yahoo.co.in

Web: www.ipcc.co.in

IPSR Solutions Ltd

Earn RHCE / RHCSS certification, in Kerala along with a boating & free accommodation. IPSR conducted more than 2000 RHCE exams with 95-100% pass rate. Our faculty panel consists of 15 Red Hat Certified Engineers.

Kochi, Kerala

Tel: +91 9447294635

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Email: training@ipsrsolutions.com
Web: www.ipsr.org

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorised training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSS, CLP, CLE, LPI - 1 & 2.

New Delhi

Mobile: 09910710143, Fax: 011-25886909
Email: info@koenig-solutions.com
Web: www.koenig-solutions.com

NACS/CIT

We are Providing Training of LINUX to Professional & Cooperate.

Meerut

Tel: 0121-2420587, Mobile: 9997526668
Email: Info@nacsglobal.com
Web: www.nacsglobal.com

NACS Infosystems (P) Ltd

NACS is a organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having testing centre of THOMSON PROMETRIC and PEARSON VUE.

Meerut

Tel: 0121-2767756, Fax: 0121-4006551
Mobile: 09897796603
Email: info@nacsglobal.com,
mohit@nacsglobal.com.
Web: www.nacsglobal.com

Netdiox Computing Systems

We are one-of-a-kind center for excellence and finishing school focusing on ground breaking technology development around distributed systems, networks, storage networks, virtualisation and fundamental algorithms optimized for various appliance.

Bangalore

Tel: 080-26640708
Mobile: 09740846885
Email: info@netdiox.com

NetMax-Technologies

Training Partner of RedHat, Cisco

Chandigarh

Tel: 0172-2608351, 3916555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com

Netzone Infotech Services Pvt Ltd

Special batches for MCSE, CCNA and RHCE on RHEL 5 with exam prep module on fully equipped labs including IBM servers, 20+ routers

and switches etc. Weekend batches are also available.

New Delhi

Tel: 011-46015674, Mobile: 9212114211
Email: info@netzoneindia.net

Plexus Software Security Systems Pvt Ltd

Plexus, incorporated in January 2003 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training; this has put Plexus in the unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training.

Chennai
Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education

RHCE & RHCSS Certifications

Jabalpur

Tel: 0761-4039376,
Mobile: 09425152831
Email: naidu.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd

Q-SOFT is in a unique position for providing technical training required to become a Linux Administration under one roof. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train on Sun Solaris, Suse Linux & Redhat Linux.

Bangalore

Tel: 080-26639207, 26544135, 22440507
Mobile: +91 9945 282834
Email: counsellors@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network

STN is one of the most acknowledged name in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.

Chandigarh

Tel: 0172-5086829
Email: stn2001@rediffmail.com
Web: stntechnologies.com

South Delhi Computer Centre

SDCC is for providing technical training courses (software, hardware,

networking, graphics) with career courses like DOEACC "O" and "A" Level and B.Sc(IT), M.Sc(IT), M.Tech(IT) from KARNATAKA STATE OPEN UNIVERSITY.

New Delhi

Tel: 011-26183327, Fax: 011-26143642
Email: southdelhicomputercentre@gmail.com,
southdelhicomputercentre@hotmail.com.
Web: www.itwhizkid.com
www.itwhizkid.org

Ssystems Quest

Making Tomorrow's professionals TODAY

Bangalore

Tel: 080-41301814
Email: directorv@ssystemsquest.com
Web: www.ssystemsquest.com

Trimax FuturePerfect

A Div of Trimax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCTS, MCITP, MCSE 03, CCNA, CCNP, Prometric Center.

Mumbai

Tel: 022-40681313, Mobile: 09987705638
Fax: 022-40681001
Email: futureperfect@trimax.in
Web: www.trimax.in

Vibrant e Technologies Ltd

Vibrant e Technologies Ltd. Is a authorised Red Hat Test and Testing Centre, has won the prestigious award " REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE 5, RHCSS etc.

Mumbai

Tel: 022-26285066/6701
Email: vibrant@vsnl.net
Web: www.vibrantcomputers.com

Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses and authorised testing center of Prometric, Vue and Red Hat.

Mumbai

Tel: 022-67669217
Email: unmesh.raote@ultramaxit.com
Web: www.ultramaxit.com



Yash Infotech

Authorized Training & Exam Center. Best Performing Center in Lucknow for RH Training and Examinations. LINUX & Open Source training institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSS, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools.

Lucknow

Tel: 0522-4043386, Fax: 0522-4043386
Email: yashinfotech.lko@gmail.com

Want to
register
your
organisation
in FOSS
Yellow
Pages For

FREE

Call: Dhiraj (Delhi) 09811206582

Somaiah (Bangalore) 09986075717

or mail: dhiraj.khare@efyindia.com

or somaiah.km@efyindia.com

*Offer for limited period.

**FOSS
Yellow
Pages**

The best place for you to buy and sell FOSS products and services



**PAY
FOR
ONE**

**GET
TWO**

Celebrating 41st Year

YOU CAN CHOOSE ANY OF THE FOLLOWING OPTIONS

**Pay for 6 months
Get for 1 year!**

**Pay for 1.5 years
Get for 3 years!**

**Pay for 2.5 years
Get for 5 years!**

MEGA SUBSCRIPTION OFFER

MAGAZINE	<input checked="" type="checkbox"/> Mark Required	Unit Price (Rs.)	Pay for 6 months Get for 1 Year	Pay for 1.5 Years Get for 3 Years	Pay for 2.5 Years Get for 5 Years
Electronics For You (Without CD)	<input type="checkbox"/>	40/-	<input type="checkbox"/> Pay Rs. 250/-	<input type="checkbox"/> Pay Rs. 730/-	<input type="checkbox"/> Pay Rs. 1210/-
Electronics For You (with CD)	<input type="checkbox"/>	60/-	<input type="checkbox"/> Pay Rs. 360/-	<input type="checkbox"/> Pay Rs. 1080/-	<input type="checkbox"/> Pay Rs. 1800/-
Electronics Bazaar	<input type="checkbox"/>	50/-	<input type="checkbox"/> Pay Rs. 310/-	<input type="checkbox"/> Pay Rs. 910/-	<input type="checkbox"/> Pay Rs. 1510/-
Linux For You (with CD & DVD)	<input type="checkbox"/>	100/-	<input type="checkbox"/> Pay Rs. 610/-	<input type="checkbox"/> Pay Rs. 1810/-	<input type="checkbox"/> Pay Rs. 3010/-
Information Technology	<input type="checkbox"/>	40/-	<input type="checkbox"/> Pay Rs. 240/-	<input type="checkbox"/> Pay Rs. 720/-	<input type="checkbox"/> Pay Rs. 1200/-
BenefIT	<input type="checkbox"/>	50/-	<input type="checkbox"/> Pay Rs. 300/-	<input type="checkbox"/> Pay Rs. 900/-	<input type="checkbox"/> Pay Rs. 1500/-
Facts For You	<input type="checkbox"/>	100/-	<input type="checkbox"/> Pay Rs. 600/-	<input type="checkbox"/> Pay Rs. 1800/-	<input type="checkbox"/> Pay Rs. 3000/-

• These rates are applicable for new subscribers as well as renewal by existing subscribers. • The rates are valid for subscribers within India only • Please allow 4-6 weeks for processing of your subscription. • Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

Hurry!

Offer Valid till 31st March 2009



A magazine for the electronics fraternity... professionals, businessmen and hobbyists




India's first Electronics Sourcing magazine




Asia's first magazine on Linux and Open Source



Information Technology Today. Tomorrow.



An IT Magazine for Business Decision Makers



A monthly on business & economic affairs with regular sections on market surveys

Subscribe Now!

ORDER FORM

Name _____ Company Name _____ Designation _____

Mailing Address _____ Distt. _____

City _____ Pin Code _____ State _____ Phone _____

Email _____ Subscription No. (for existing subscribers only) _____

I would like to subscribe to the above (✓)marked magazine(s), for the period _____ . Please find enclosed a sum of Rs _____ by DD/MO/crossed cheque*bearing the No. _____ dt. _____ in favour of **EFY Enterprises Pvt Ltd, payable at Delhi.** (*Please add Rs 50 on non-metro cheque)

or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs. _____

against my credit card No. _____

Date of Birth ____/____/____ (dd/mm/yy) Card Expiry Date ____/____/____ (mm/yyyy)



Signature (as on the card)



Send this filled-in form or its photocopy to :

EFY Enterprises Pvt Ltd: D-87/1, Okhla Industrial Area, Phase I, New Delhi 110 020

Ph: 011-26810601-03; Fax: 011-26817563, E-mail: info@efyindia.com website: www.efyindia.com

An Exclusive B2B Exhibition - With a Senior Level Conference

Connect 2009

4th Information & Communications
Technology Exhibition & Conference

5 - 7 May 2009

Karachi Expo Centre

www.connectit.com.pk



ORGANISED BY



Pegasus Consultancy (Pvt) Ltd

2nd Floor, Business Centre,
Mumtaz Hassan Road, Karachi - Pakistan
Tel : +(92 21) 111 734 266 (PEGCON)
Fax : +(92 21) 241 0723

E-mail : info@connectit.com.pk, URL : www.pegasus.com.pk

ACTIVITY PARTNER



National Response Centre
for Cyber Crime (FIA)

PARTNER SHOW



SUPPORTED BY



**SPECIAL
OFFER**

SUBSCRIBE



Lifeline of Cricket
Cricket
today
www.crickettoday.com



Get exiting Gifts



TIMEX
USA SINCE 1954

Rs 1500/-



Rs 749/-

+



Rs 395/-

FREE
It's a profit game


☐

Yes, I would like to subscribe Diamond Cricket Today english for the following term

**3 Year (36 Issues). I pay Rs. 1800/-
get Timex watch worth Rs. 1500* +
an IPL book worth Rs. 395/- FREE!**

☐

**2 Year (12 Issues). I pay Rs. 1200/-
get Allen Solly tie worth Rs. 749* +
an IPL book worth Rs. 395/- FREE!**



Name: Mr./Ms.....Date of birth □□/□□/□□

Address.....

.....PIN □□□□□□

Occupation:.....

Ph:.....E-mail:.....

Please find enclosed () Cheque/DD No.....Dated.....

for Rs. 1800 ☐ Rs. 1200 ☐ favouring Diamond Magazines Pvt. Ltd.

Please charge it to my ☐ Visa ☐ Master Card No. □ □ □ □ □ □ □ □ □ □ □ □ □ □

Card expiry date □□/□□ Verification Code □ □ □ (This no. is the last three digits of the number printed at the back side of your credit card)

Date.....Signature.....

Please fill in this order form and mail it to:


**Cricket Today, Diamond Magazines Pvt. Ltd., X-30 Okhla Industrial Area, Phase-II,
New Delhi-20 Ph. 41611861 Fax: 41611866 Email: crickettoday@dpb.in**

*Please note: The gift may vary from the one shown above.



Lifeline of Cricket
Cricket
today

A solution **FLEXIBLE** enough
to **ADAPT** to your requirements



UP TO 2GB/s
UP TO 576 TB
SRP Support

Opslag FS2

Unified NAS & SAN

Flexible Host Interface

Supports Gigabit Ethernet, 10GB Ethernet,
SDR/DDR/QDR (10Gbps/20Gbps/40Gbps) as individual connectivity or in combination

Scalable Storage

Supports up to 576 TB from few TB's using Tyrone expandable solutions

Redundancy/ Fault-tolerance

Power-supply, Host interface failover, RAID 6 (various level of RAID supported), Multiple Snapshots with scheduling, Volume/
Share Replication, Remote Replication (Block level replication to a 2nd unit with failover/high availability)

Supports Virtual Tape Library

Block-Storage/SAN Protocols

iSCSI target, Fibre-Channel target, SRP target

Supported Clients

Windows, Linux, FreeBSD*, Solaris*, Unix* (*no SRP support)